Numerical Weather Prediction

STEPS: A probabilistic precipitation forecasting scheme which merges an extrapolation nowcast with downscaled NWP

Forecasting Research Technical Report No. 433

Neill Bowler, Clive Pierce and Alan Seed[1]

[1] Weather forecasting group, Bureau of Meteorology Research Centre, Melbourne, Australia

email: nwp_publications@metoffice.com

©Crown Copyright

STEPS: A probabilistic precipitation forecasting scheme which merges an extrapolation nowcast with downscaled NWP

Neill Bowler[1*] Clive Pierce[1†]

Alan Seed[2‡]

[1]Met Office, JCHMR, Maclean Building, Crowmarsh Gifford, Wallingford, Oxon, OX10 8BB.

[2]Cooperative Research Centre for Catchment Hydrology, Weather Forecasting Group, Bureau of Meteorology Research Centre, Melbourne, Australia.

March 31, 2004

Abstract

An ensemble-based probabilistic precipitation forecasting scheme has been developed that blends an extrapolation nowcast with a downscaled NWP forecast, known as STEPS: Short-Term Ensemble Prediction System. The uncertainties in the motion *and* evolution of radar inferred precipitation fields are quantified, and the evolutionary uncertainty is shown to be the more important. The use of ensembles allows the scheme to be used for applications that require forecasts of the probability density function of areal and temporal averages of precipitation, such as fluvial flood forecasting - a capability that has not been provided by previous probabilistic precipitation nowcast schemes. The output from an NWP forecast model is downscaled so that the small scales not represented accurately by the model are injected into the forecast using stochastic noise. This allows the scheme to better represent the distribution of precipitation rate at spatial scales finer than those adequately resolved by operational NWP. The performance of the scheme has been assessed over the first 10 days of the month of March 2003. Performance evaluation statistics show that the scheme possesses predictive skill at lead times in excess of six hours.

*Electronic address: Neill.Bowler@metoffice.com

†Electronic address: Clive.Pierce@metoffice.com

‡Electronic address: A.Seed@bom.gov.au

Executive Summary

Although Numerical Weather Prediction (NWP) models are the primary source of weather forecast guidance in the Met Office and other National Weather Services, currently, they are not suitable for nowcasting applications, because they are too expensive to run with the required update cycle (every 15 minutes), and are unable to predict the distribution of precipitation at the required spatial ($\leq$ 5 km) and temporal ($\leq$ 15 minute) resolutions with sufficient accuracy. By contrast, nowcasts based upon extrapolation of radar inferred precipitation, whilst being of superior accuracy (at lead times of less than $\sim$ 3 hours), lose their value very rapidly. Thus, an optimal nowcast is obtained by exploiting these two, complementary approaches in combination.

In the Gandolf system the primary data inputs are C-band radar data obtained from the UK weather radar network, and the latest precipitation forecast produced by the Met Office's meso-scale model. An extrapolation nowcast obtained by advecting the latest radar inferred instantaneous rain rate distribution (hereafter referred to as a *rain analysis*) with a velocity field diagnosed from a temporal sequence of recent rain analyses is blended (a process referred to as *merging*) with time synchronous forecasts of rain rate from the most recent run of the meso-scale model. Since the horizontal resolution of the latter component ($\sim$ 12 km) is less than that of the extrapolation component (2 km), and the merging process gives increasing weight to the NWP component with advancing lead time, the resultant nowcast becomes progressively more smooth. This process broadly reflects the loss of deterministic predictive skill in the extrapolation nowcast at the smallest scales. However, this loss of information has a significant deleterious impact on the value of the nowcast in fluvial flow and flood prediction applications.

An alternative method of accounting for uncertainties in the nowcast rain rate distribution is to adopt an ensemble approach in which the statistical properties of the radar inferred precipitation are conserved. Stochastic noise with the appropriate spatial and temporal correlation structures can be injected into a nowcast to replace extrapolated precipitation features in a manner that is consistent with the loss of forecast skill with advancing lead time. Since this stochastic noise is sampled from a space-time probability density function of instantaneous rain rate, multiple forecast scenarios must be generated in order to represent the associated uncertainty. This approach is advantageous from a hydrological forecasting perspective in that it offers a means of estimating a PDF of fluvial flows, and hence of predicting the probability of a fluvial flood.

A new, probabilistic quantitative precipitation nowcast (PQPN) scheme has been developed based upon the ensemble approach outlined above. This exploits the modelling framework previously adopted in the Nimrod and Gandolf QPN schemes. Uncertainties in the predicted motion and Lagrangian evolution of radar inferred precipitation fields are quantified. The uncertainty in the field motion is modelled by adding a perturbation to the advection velocity field. This perturbation is consistent with the measured evolution of the velocity field over time. The evolution of instantaneous rain rate is modelled in the frequency domain by decomposing a rain analysis into a cascade of features representing the field variability on a hierarchy of

spatial scales. This framework provides a means of modelling the scale dependent temporal evolution of precipitation over time. Extrapolated precipitation features are replaced with scale filtered, temporally correlated stochastic noise, at a rate that is consistent with the observed temporal persistence of these precipitation features. Merging with NWP precipitation forecasts is achieved in such a way that stochastic noise is injected at those scales inadequately represented by the meso-scale model.

1 Introduction & Literature Review

The topic of probabilistic quantitative precipitation forecasting (PQPF) has previously received attention in the form of model output statistics [2] (MOS) and ensemble prediction systems [22, 36] (EPSs) using NWP models. There has been comparatively little work on PQPF using nowcasting methods (which will be abbreviated to PQPN). This work combines a probabilistic extrapolation nowcast with an NWP downscaling scheme to produce probabilistic forecasts of precipitation for lead times up to 6 hours (the maximum lead time is only limited by computer resources and the availability of NWP forecasts).

Deterministic nowcasting models have been shown to exhibit useful predictive skill at lead times of many hours [9]. Nowcasting relies on the use of up to date observations of weather and simple forecasting methodologies to produce precipitation forecasts with high spatial ($\leq$ 5 km) and temporal ($<$ 1 hour) resolutions. Optimal performance is achieved by blending these nowcasts with NWP model forecasts [10, 11] in such a way that the "merged" product is more skillful than either component at lead times between approximately 1 and 5 hours. Since a probabilistic forecast is produced here the forecasts are expected to provide useful guidance at all lead times.

For many users of precipitation forecasts, the output available from current NWP models does not have a sufficiently high spatial resolution. In hydrological modelling many important river catchments are only a few km across. However, the Met Office's UK meso-scale model is run at a grid length of approximately 12km, and can only represent systems which are on a scale much larger than this. A convenient rule of thumb is that NWP models do not accurately represent features on a scale of less than 5 times the grid length of the model [12]. Clearly an effective resolution of 60km is not sufficient for hydrological modelling. Since the NWP forecast does not contain sufficient detail at the smallest scales it is possible to insert that detail using a statistical model. The forecasts produced are therefore intrinsically probabilistic and the details is generated randomly and cannot be relied upon.

It has long been noted that rainfall patterns show scaling behaviour consistent with multi-fractal models [8, 21, 37]. This scaling may be exploited in the downscaling of NWP model forecasts of rainfall using multiplicative cascades with the rainfall at a point being given by the product of a cascade of values containing information on a series of spatial scales. The simplest application of this is seen in the wavelet transform approach [14, 15]. The essence of this approach is to downscale the coarse-scale rainfall by splitting each grid box into 4 smaller boxes and setting the rain rate in the smaller box to a multiple of a randomly chosen number and the rain rate at the coarser scale (see fig. 1). The four random numbers are constrained so that their average is one (and the total rain amount is preserved). Repetition of this process allows a rain field to be downscaled to as fine resolution as is desired.

Spatio-temporal models have been proposed [18, 24] which allow a consistent series of downscaled rain rate fields to be produced. The String of Beads model [24] uses a diagnosed advection velocity with an auto-regressive model to account for the temporal evolution of the rain field. Other authors [18] have considered the rain /

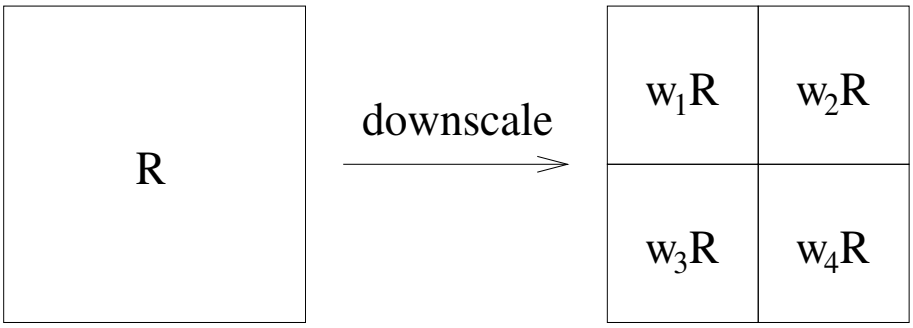

Figure 1: In a simple downscaling approach the coarse-scale grid box is split into 4 boxes, with the rain rate in each of the smaller boxes to a randomly chosen number (≥ 0) times the coarse scale rain rate. The average rain rate at the smaller scale may be constrained to equal the coarse-scale rain rate.

no-rain process separately from the variability of the rain using a Bayesian method from the field of image analysis. The downscaling is updated at each new time step without reference to an advection model. The linkage between meteorological parameters and parameters controlling the cascade has been considered [25, 26] with results indicating that the amount of convectively available potential energy (CAPE) has a significant influence on the structure of the observed rainfall.

We know of three PQPN schemes currently in existence, all of which rely on a smearing of the extrapolation forecast. In the case of the Raincast system [30, 31], the PDF of an extrapolated rain rate field is estimated from a weighted combination of rain rate values taken from a deterministic forecast surrounding the grid point of interest. The weight given to a specific grid point in this computation is taken from a Gaussian distribution whose width is allowed to increase linearly with lead time. The extrapolation algorithm employed is the COTREC scheme [20, 17].

A probabilistic precipitation nowcasting scheme has been used operationally within the Met Office for a number of years [34]. Essentially, this is a simplified version of the Raincast scheme. The probability of precipitation is estimated at each grid point of interest by considering the fraction of precipitating pixels in a box surrounding the chosen grid point. The size of this box increases linearly with lead time.

The scheme proposed by Andersson & Ivarsson [1] is the earliest of the probabilistic schemes known to us. Again the philosophy is similar. A precipitation nowcast is generated by an advection scheme, in this case based upon wind forecasts from an NWP model. A parameterised PDF is estimated from the extrapolated distribution of rain rates within a circle surrounding a point upwind of the location for which a probabilistic forecast is required.

All of the above schemes approximate the PDF of the rain rate from a single forecast scenario by assuming that the uncertainty in the rain rate distribution derives solely from errors in the motion of the precipitation field. Clearly, this approach is deficient because no account is taken of growth and decay, which is believed to

be the major failing of nowcasting systems [33]. These PQPN schemes have been unable to provide adequate guidance where an accumulation (over time or over an area) is required since the shape and motion of a rain area may affect the total. Hence these schemes are unsuitable for many applications, such as hydrological modelling. For this reason, we have developed a scheme that produces an ensemble of possible future rain rates whilst attempting to quantify uncertainties in the evolution of precipitation fields, including those normally attributable to growth and decay, and advection. Furthermore, by combining these nowcasts with downscaled NWP forecasts the useful range of the resultant forecasts can in theory be extended to a number of days.

The data used in this study are the radar rain analyses for the Gandolf [3, 4, 27] domain and the Met Office meso-scale NWP forecasts of rain rate. The domain covers an area of 1150 km by 1150 km of the UK National Grid, with top left corner at -290 km E, 910 km N. This domain has been trimmed to 1024 km by 1024 km by removing the most southerly and easterly 126 km. The radar data over the UK are subject to a number of quality control and correction algorithms [13]. From this a surface rain rate is estimated, which will henceforth be known as a rain analysis. For the purposes of this work the rain analysis will be assumed to be error-free. This assumption is necessary due to the current absence of any credible estimates of the errors present in the rain analysis.

Both the rain analyses and the NWP forecasts are available as surface rain rate values in units of mm/h. For the S-PROG decomposition (see below) it is necessary to apply a logarithmic transformation to the rain rate values. This transformation follows

$$Z = 10 \log_{10} \left(200 R^{1.6} \right) \tag{1}$$

where R is the estimated surface rain rate. For areas where no precipitation is detected or the above formula gives a value less than 10, Z is set to -6. Values for Z of -6 and 10 corresponds to rain rates of 0.025 mm/h and 0.068 mm/h, respectively. The choice of -6 and 10 is arbitrary and dictated by which values produce the best results for the forecasting scheme. This transformation is similar to the conversion from rain rate (in mm/h) to radar reflectivity (dBZ). However, the choice of transformation is determined the the requirements of the S-PROG scheme alone.

The estimation of parameters relating to the cascade has been performed on rain analyses and NWP forecasts for April 2003. Only times when more than 5% of the radar area contained precipitation have been used. 5% is a compromise between considering only cases with significant amounts of rain and not limiting the set of available data too severely. The work on advection velocities has been performed on data from November 2002. The forecasts have been run on data from March 2003 as it is good practise to avoid running forecasts on the data set used for parameter estimation.

2 Development of Precipitation Patterns

2.1 Spectral Prognosis (S-PROG)

In recent years substantial theoretical and empirical evidence has been produced to suggest that precipitation fields exhibit spatio-temporal properties that are consistent with their representation by multi-fractal models [8, 21, 37]. These models assume that similar features exist in precipitation fields on a continuum of scales from the large (10^4 m) to the very small ($<$ 10 m). Observations show that the temporal persistence of such features is broadly proportional to their size. These fundamental spatio-temporal properties of the precipitation field underpin the design of the Spectral-Prognosis (S-PROG) model [32] and form the basis for the modified version of S-PROG described below.

A field of transformed rain rate (see eq. 1) with dimensions of L by L pixels may be decomposed into an additive cascade

$$Z_{i,j}(t) = \sum_{k=1}^{n} X_{k,i,j}(t) \text{ where } i = 1, \ldots, L, j = 1, \ldots, L, L = 2^n \qquad (2)$$

where the k_{th} field in this cascade, $X_k(t)$, represents the variability in the original field with frequencies ω_k within the range $\frac{2^{k-1}}{L} < \omega_k < \frac{2^k}{L}$ pixel^{-1} at time t. Since Z is proportional to the logarithm of the rain rate, the cascade described here is multiplicative when expressed in terms of the rain rate (mm/h).

The transformed rain rate field, Z is translated into the spectral domain by means of a Fast Fourier Transform (FFT). Each level $X_k(t)$ of the cascade is calculated using a band-pass filter (based on a Gaussian window) to pass the appropriate frequencies. An inverse transform is performed to return each Fourier component of the original field back into the spatial domain. Finally, the cascade levels are normalised using

$$Y_{k,i,j}(t) = \frac{X_{k,i,j}(t) - \mu_k(t)}{\sigma_{k,i,j}(t)} \qquad (3)$$

where $\mu_k(t)$ and $\sigma_{k,i,j}(t)$ are the mean and standard deviation of the k^{th} level respectively. $\sigma_{k,i,j}(t)$ is allowed to vary in space, but this variation is constrained to be much slower than the variation in $X_{k,i,j}(t)$. This is achieved by calculating the value of the variance of $X_{k,i,j}(t)$ over a series of boxes of size $\min\left(L2^{-k+4}, L\right)$, and linearly interpolating the value of the standard deviation at points not lying at the centre of a box to avoid sharp jumps in $\sigma_{k,i,j}(t)$. If the variance in one of these boxes is found to be less than $1/100^{th}$ the variance found over the whole domain, then the variance of that box is set to the large-scale variance. This avoids generating only smooth features where there are currently rain-free areas. The variance at small scales is typically within a small factor of the large scale variance, so the threshold is unlikely to cause problems. The $Y_{k,i,j}$ are referred to as the component "levels" of the cascade and represent the Fourier components transformed into the spatial domain. These are normalised to have zero mean and unit variance. Figure 2 shows a field of radar reflectivity (similar to the transformed rain rate, see eq. 1). The first three component levels, $Y_{k,i,j}$, of the cascade are reproduced in figure 3.

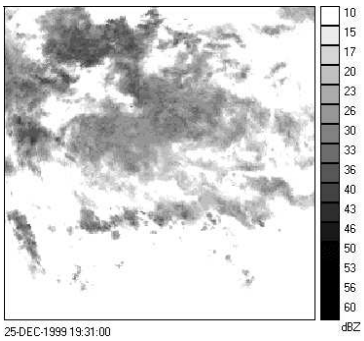

Figure 2: A field of radar reflectivity. The Melbourne (Laverton) radar is at the centre; the image is a 256 km x 256 km field with 1 km resolution data.

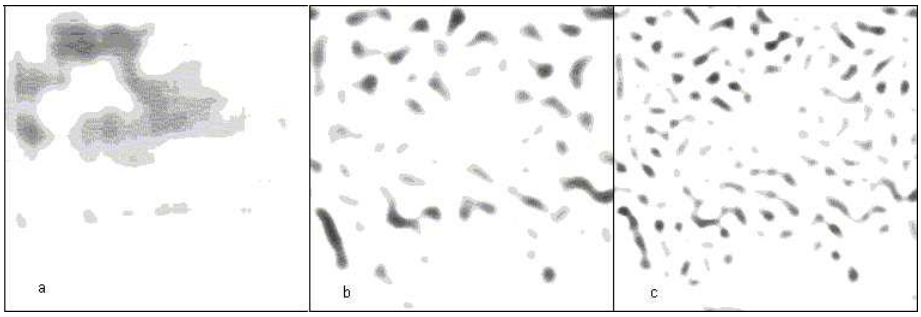

Figure 3: The first three levels in the cascade of the field in figure 2 a) 128 - 256 km, b) 64-128 km, c) 32 - 64 km wavelengths. Note that detail is seen in the smaller levels away from the precipitation area due to "ringing".

The decomposition used here is based on the Fourier transform, but one may consider other decomposition algorithms, such as wavelet transforms [16]. However, the undesirable "ringing" (Gibbs effect) that is seen in figure 3 is unavoidable for all decompositions. An approximately Gaussian window is used to perform the filtering which offers a compromise between exact localisation in spatial scale and avoiding the ringing. The differences between different decompositions are small and there is no clear reason for choosing one over another.

2.2 Extrapolation, Noise & NWP Cascades

The production of a merged probabilistic forecast is performed by the maintenance of 3 cascades; one for the extrapolation forecast, one for the NWP forecast and one for the noise.

As the forecast progresses the extrapolation component will become less skillful, especially at the smallest scales. It was proposed by Seed [32] that this loss of skill can be modelled using a hierarchy of second-order auto-regressive models (AR-2). These allow precipitation features to be smoothed out at a rate which is consistent with their observed temporal persistence. This smoothing is achieved by decreasing the power in the lowest levels of the cascade in accordance with the observed auto-correlation of features at that scale. The power lost at the small scales can be replaced by noise, thereby permitting the uncertainty at these scales to be modelled effectively. A cascade of noise may be created using the same Fourier filtering method described above for the rain analyses. This will allow the injection of noise at the appropriate scales. Thus, a probabilistic nowcast may be thought of as a merging of an extrapolation and a noise cascade.

In order to produce nowcasts with a useful range of at least six hours it is necessary to merge these nowcasts with NWP forecasts. This merging will allow the resultant precipitation nowcasts to reflect the influence of the large-scale dynamical evolution of the atmosphere on the precipitation field. The NWP forecast may be decomposed into a cascade in the same way as discussed above. This cascade is expected to contain less variability at the small scales than the cascade for the radar analyses. Therefore, by merging the NWP cascade with a noise cascade (which will provide the extra variability at the small scales) the NWP forecast may be appropriately downscaled to be statistically equivalent to the rain analysis.

The production of a probabilistic forecast which merges the extrapolation forecast and NWP forecast may be achieved by merging the extrapolation, noise and NWP cascades according to

$$Y_{k,i,j}(t+t_l) = w_k^e(t+t_l)Y_{k,i,j}^e(t+t_l)+w_k^m(t+t_l)Y_{k,i,j}^m(t+t_l)+w_k^n(t+t_l)Y_{k,i,j}^n(t+t_l) \quad (4)$$

where Y^e, Y^m and Y^n denote the extrapolation, NWP model and noise cascades, respectively. These cascades are given weights w^e, w^m and w^n. The forecast is made at a observation time (data time) of t, and t_l is the lead time. The weighting used is described in section 2.6. To allow the merging to be performed it is necessary to maintain a noise cascade with the correct temporal correlations, which is discussed in the next section.

2.3 Lagrangian Evolution of the Noise & Extrapolation Cascades

As previously discussed, other PQPN schemes have attempted to model uncertainties in advection velocity alone. Here we describe our model for the uncertainties arising from growth and decay of the precipitation field. The evolution of the field at each level in the extrapolation cascade is modelled using a second order Auto-Regressive (AR-2) process [5]

$$Y^e_{k,i,j}(t+t_l) = \phi_{k,1}(t)Y^e_{k,i,j}(t+t_l-\Delta t) + \phi_{k,2}(t)Y^e_{k,i,j}(t+t_l-2\Delta t) \quad (5)$$

where Δt is the time between forecast steps. The noise term has been omitted from this formulation and the cascade is re-normalised at each step to ensure no loss of power. This is similar to the original S-PROG approach [32] and is chosen since the noise will be maintained in a separate cascade.

The evolution of the noise cascade is driven by the same AR-2 model according to

$$\begin{aligned} Y^n_{k,i,j}(t+t_l) \quad = \quad & \phi_{k,1}(t)Y^n_{k,i,j}(t+t_l-\Delta t) + \phi_{k,2}(t)Y^n_{k,i,j}(t+t_l-2\Delta t) + \\ & \phi_{k,0}(t)\epsilon_{k,i,j}(t+t_l) \end{aligned} \quad (6)$$

where $\epsilon_{k,i,j}$ is an independent cascade of noise generated for each time step. Thus a cascade of temporally correlated noise is maintained for the range of the forecast. Both this cascade and the extrapolation cascade are advected using the current estimate of the velocity field (see section 2.7). The parameters $\phi_{k,0}$, $\phi_{k,1}$ and $\phi_{k,2}$ control the rate of temporal evolution at each scale, and are determined from the sequence of rain rate analyses. The estimation of $\phi_{k,1}$ and $\phi_{k,2}$ is dealt with below. The noise fields are chosen to be Gaussian [21], $N(0,(\sigma^n_k)^2)$, and $\phi_{k,0}$ is given by [29]

$$\phi_{k,0}(t) = \sqrt{\frac{1+\phi_{k,2}(t)}{1-\phi_{k,2}(t)}\left[\left(1-\phi_{k,2}(t)\right)^2 - \left(\phi_{k,1}(t)\right)^2\right]} \quad (7)$$

which ensures that $Y^n_{k,i,j}(t+t_l)$ has the same normalisation as $Y^n_{k,i,j}(t)$ (mean zero and variance $(\sigma^n_k)^2$).

The levels in the noise cascade are not chosen to be normalised to standard deviation 1, but instead σ^n_k. It is essential that the noise has the same statistical properties as the precipitation field from the rain analysis, which is not the case if the levels in the noise cascade have variance 1. When calculating the rain rate from a cascade large areas will be diagnosed as being free from precipitation since the transformation of the rain rate involves a threshold. Hence some of the variability in the cascade will not be represented in the rain rate field. Furthermore, since the Z values between -6 and 10 are eliminated in the conversion there will be increased variability in the small-scales as the cascade attempts to capture large jumps in the transformed rain rate. These choices of thresholds of -6 and 10 in the transformed rain rate are entirely arbitrary and have been chosen to optimise performance of the forecast scheme.

The appropriate values for σ_k^n may be measured by generating a cascade of noise, re-composing this cascade and calculating a rain rate field. This field may be transformed back in Z values and decomposed in the usual manner. Comparison of the cascade parameters for the original cascade and the cascade which has been through conversion to rain rate will yield values for σ_k^n. From this analysis it has been found that

$$\sigma_k^n = \frac{1}{0.75 + 0.09k} \tag{8}$$

is a suitable fit, which indicates that the small-scale variance tends to be over-estimated, whereas the large-scale variance is under-estimated (in our computations, 8 levels were used).

2.4 Estimation of the Auto-Regression parameters

The values of ϕ (parameters of the AR-2 model) are estimated using a temporal sequence of three rain analyses spanning a half hour period. This involves deriving an advection velocity field and computing the correlation coefficients γ using

$$\gamma_{k,l} = \frac{\sum_{i=1}^{N}\sum_{j=1}^{N} Y^e_{k,i,j}(t) Y^e_{k,i,j}(t - (l\Delta t))}{\sqrt{D_{k,0} D_{k,l}}} \tag{9}$$

where

$$D_{k,l'} = \sum_{i=1}^{N}\sum_{j=1}^{N} Y^e_{k,i,j}(t - (l'\Delta t)) Y^e_{k,i,j}(t - (l'\Delta t)), \tag{10}$$

and $l = 1, 2$ is the lag of the correlation coefficient, in this case measured in time steps of Δt minutes. $Y^e_{k,i,j}(t - (l\Delta t))$ has been advected to the current time, t, so that the Lagrangian correlation is measured, and N is the size of the rain analysis. The estimated values for $\gamma_{k,l}$ are temporally smoothed according to

$$\gamma_{k,l}(t) = 1 - \sqrt{(1 - \bar{\gamma}_{k,l}(t))(1 - \gamma_{k,l}(t - \Delta t))} \tag{11}$$

where $\gamma_{k,l}(t - \Delta t)$ is the time-smoothed value at the previous time step and $\bar{\gamma}_{k,l}(t)$ is the unsmoothed estimate at the current time. Using the geometric mean of the difference from perfect correlation gives suitable behaviour of the average when $\gamma_{k,l}$ is close to 1. The auto-correlation coefficients are allowed to regress towards their climatological values (long-term averages) as a forecast progresses according to the observed autocorrelation of the measured values which is

$$\gamma_{k,l}(t + t_l) = 1 - (1 - \gamma_{k,l}(t))^{q_{r,k}} (1 - \hat{\gamma}_{k,l})^{1 - q_{r,k}} \tag{12}$$

where $q_{r,k}$ controls the speed of the regression and is

$$q_{r,k} = e^{-t_l / L_{r,k}} \tag{13}$$

Level	Spatial scales (km)	$L_{r,k}$
0	256-512-1024	165
1	128-256-512	125
2	64-128-256	145
3	32-64-128	170
4	16-32-64	305
5	8-16-32	375
6	4-8-16	270
7	Residual	60

Table 1: The regression parameter for the auto-correlation coefficients.

where t_l is the lead time in minutes and the values of $L_{r,k}$ are chosen to match the observed temporal persistence of the autocorrelation coefficients and are given in Table 1. Measurements of the average lag 1 & 2 auto-correlation coefficients are shown in Table 2. The data have been measured over April 2003 and averaged using the geometric mean of the deviation from perfect correlation. The AR-2 parameters are then calculated using [5]

$$\phi_{k,1}(t+t_l) = \frac{\gamma_{k,1}(t+t_l)(1-\gamma_{k,2}(t+t_l))}{1-(\gamma_{k,1}(t+t_l))^2} \tag{14}$$

$$\phi_{k,2}(t+t_l) = \frac{\gamma_{k,2}(t+t_l)-(\gamma_{k,1}(t+t_l))^2}{1-(\gamma_{k,1}(t+t_l))^2}. \tag{15}$$

We choose to restrict the parameters of the AR-2 model using the following conditions [5]

$$(\phi_{k,1}(t+t_l))^2 + 4\phi_{k,2}(t+t_l) \geq 0 \tag{16}$$

and

$$\phi_{k,1}(t+t_l) \geq 0. \tag{17}$$

Together these conditions ensure that the solutions produced by the model are non-oscillatory (the auto-correlation is always positive). Therefore, we are choosing to ignore solutions where an area of rain is regularly appearing and disappearing. Rather than restrict the values of the AR-2 parameters we restrict the auto-correlation coefficients. Therefore, we insist that

$$\gamma_{k,2}(t+t_l) \geq \frac{3\gamma_{k,1}^2(t+t_l) - 2 + 2(1-\gamma_{k,1}^2(t+t_l))^{\frac{3}{2}}}{\gamma_{k,1}^2(t+t_l)} \tag{18}$$

and

$$\gamma_{k,1} \geq 0. \tag{19}$$

It should be noted that the climatological values of the auto-correlation coefficients (Table 2) do not all meet these criteria.

Level	Spatial scales (km)	Lag 1	+/-	Lag 2	+/-
0	256-512-1024	0.99805	0.00007	0.9933	0.0002
1	128-256-512	0.9925	0.0002	0.9752	0.0005
2	64-128-256	0.9776	0.0004	0.923	0.001
3	32-64-128	0.9297	0.0009	0.750	0.003
4	16-32-64	0.796	0.002	0.367	0.005
5	8-16-32	0.482	0.005	0.069	0.003
6	4-8-16	0.079	0.004	0.0018	0.0009
7	Residual	0.0006	0.0008	0.0014	0.0006

Table 2: The climatological values of the lag 1 and lag 2 auto-correlation coefficients.

2.5 Model variability in cascade levels

Since precipitation forecasts from numerical weather prediction models have been shown to be too smooth [12], it is expected that extra variability will have to be injected into the NWP forecast at scales below those adequately resolved. One may decompose an NWP forecast in the same way as is proposed for the rain analyses, and compare the variability of the analysis and a concurrent forecast at each level in the two cascades. For this comparison $\sigma_{k,i,j} = \sigma_k$ is constrained to a single value for the entire domain, and both the forecast and analysis are set to the "no-data" value wherever either shows that value. The comparisons are made for every NWP forecast in the month of April 2003 when more than 5% of the analysis contained precipitation, and the ratio between the variability at each level is averaged over this month. The lead time of the NWP forecasts used varied in the range 4h to 18h. The following was calculated

$$\hat{f}_k = \left\langle \frac{\sigma_k^r}{\sigma_k^m} \right\rangle \tag{20}$$

where r denotes the variability (standard deviation) diagnosed from the radar analysis and m denotes the variability of concurrent NWP forecast.

The average values of the ratio between the variability in the cascade levels is shown in Table 3. The NWP forecasts contain too little variability at the smallest scales, as expected. However, they also appear to contain too much variability at the largest scales. This could be a symptom of either the NWP model forecasting precipitation to occur over too wide an area or the radar being unable to observe large areas of light rain / drizzle.

The ratio may be estimated using the latest NWP forecast and these values of the ratio will be used to adjust $\sigma_{k,i,j}^m$ so that an unbiased estimate of the variability in the rain analysis is provided by the NWP forecast. These latest estimates of the ratio are allowed to regress towards $\hat{f}_k$ as the forecast progresses. The value used at a lead time t_l (in minutes) is given by

$$f_k(t+t_l) = e^{-t_l/200}\left(2 - e^{-t_l/320}\right) f_k(t) + (1 - e^{-t_l/200}\left(2 - e^{-t_l/320}\right))\hat{f}_k \tag{21}$$

Level	Range of spatial scales (km)	Climatological ratio	+/-
0	256-512-1024	0.675	0.006
1	128-256-512	0.711	0.006
2	64-128-256	0.824	0.004
3	32-64-128	0.950	0.005
4	16-32-64	1.114	0.005
5	8-16-32	1.255	0.006
6	4-8-16	1.368	0.006
7	Residual	2.335	0.009

Table 3: The climatological values of the ratio between the variability in the rain analysis cascade the NWP forecast cascade. The ratio is greater than one at the smallest scales, indicating that the rain analysis contains more structure on these scales than does the NWP forecast. At the larger scales the ratio is less than one, indicating that the radar has less structure on these scales than the NWP forecast.

where $f_k(t)$ is the latest estimate of the ration and the merging takes place over a time scale consistent with the measured auto-correlation of the estimates of the ratios.

2.6 Merging of the Three Cascades

As discussed above, the production of a probabilistic forecast which uses extrapolation and NWP forecasts may be achieved by merging an extrapolation, a noise and an NWP cascade. This section considers the weights that need to be applied to each of those cascades in the merging process.

The expected skill of the extrapolation nowcast may be found from the AR-2 model used to update the extrapolation and noise cascade. The autocorrelation of the cascades controlled by an AR-2 model is given by [5]

$$r_k^e(t+t_l) = \phi_{k,1} r_k^e(t+t_l-\Delta t) + \phi_{k,2} r_k^e(t+t_l-2\Delta t) \tag{22}$$

with the initial values of $r_k^e(t) = 1$ and $r_k^e(t+\Delta t) = \frac{\phi_{k,1}}{1-\phi_{k,2}}$. $r_k^e(t+t_l)$ is taken as the expected skill of the extrapolation component.

The expected skill of the NWP forecast is taken from the correlation of the cascade derived from the latest rain analysis with a time-synchronous cascade derived from an latest NWP forecast. The correlation coefficient, $r_k^m(t)$, between each level in the cascade from the most recent rain analysis and the corresponding NWP forecast is taken as the estimate of the current skill of the NWP forecast. During the forecast period the current skill is allowed to regress towards the climatological value of the skill according to

$$r_k^m(t+t_l) = q_k^m r_k^m(t) + (1-q_k^m)\hat{r}_k^m \tag{23}$$

Level	Range of spatial scales (km)	$L^m_{1,k}$ (minutes)	$L^m_{2,k}$ (minutes)
0	256-512-1024	130	155
1	128-256-512	165	220
2	64-128-256	120	200
3	32-64-128	55	75
4	16-32-64	50	100000
5	8-16-32	15	100000
6	4-8-16	15	100000
7	Residual	10	100000

Table 4: The parameters of the fit to the measured auto-correlation of the estimates of the NWP model skill.

Level	Range of spatial scales (km)	Climatological skill	+/-
0	256-512-1024	0.848	0.003
1	128-256-512	0.537	0.005
2	64-128-256	0.237	0.004
3	32-64-128	0.065	0.002
4	16-32-64	0.020	0.001
5	8-16-32	0.0044	0.0006
6	4-8-16	0.0052	0.0002
7	Residual	0.0040	0.0003

Table 5: The climatological values of the skill of the NWP model forecast (correlation between the NWP forecast cascade and rain analysis cascade). Even at the largest scales the NWP forecast only explains 71% of the variance in the observations. This may be due to displacement errors in the NWP forecast [7].

where

$$q^m_k = e^{-t_l/L^m_{1,k}} \left(2 - e^{-t_l/L^m_{2,k}}\right) \tag{24}$$

and t_l is the lead time of the forecast (in minutes) and the coefficients $L^m_{1,k}$ and $L^m_{2,k}$ are the measured auto-correlation of the estimates of the NWP forecast skill shown in Table 4. The climatological values of the skill, $\hat{r}^m_k$, are given in Table 5. Both of these are estimated from rain analyses and all NWP forecasts made in April 2003 when more than 5% of the rain analysis contained precipitation. The available forecasts have lead times in the range 4h to 18h. The climatological values of the NWP forecast skill may seem surprisingly low. However, if the NWP forecast suffers from a significant displacement error, then the correlation could easily become negative. Such a scenario may occur when the large-scale forcing of the precipitation is weak.

Given an estimate for the skill of the individual forecast cascades the appropriate forecast given by eq. 4. The weights in this equation are calculated from the

expected skill of the extrapolation and NWP model forecasts according to

$$w_k^e(t+t_l) = r_k^e(t+t_l)\sqrt{\frac{\lambda_k^e(t+t_l)}{\lambda_k^e(t+t_l)+\lambda_k^m(t+t_l)}} \quad (25)$$

$$w_k^m(t+t_l) = r_k^m(t+t_l)\sqrt{\frac{\lambda_k^m(t+t_l)}{\lambda_k^e(t+t_l)+\lambda_k^m(t+t_l)}} \quad (26)$$

$$w_k^n(t+t_l) = \sqrt{1-(w_k^m(t+t_l))^2-(w_k^e(t+t_l))^2} \quad (27)$$

where

$$\lambda_k^e(t+t_l) = \frac{(r_k^e(t+t_l))^2}{1-(r_k^e(t+t_l))^2} \quad (28)$$

$$\lambda_k^m(t+t_l) = \frac{(r_k^m(t+t_l))^2}{1-(r_k^m(t+t_l))^2}. \quad (29)$$

The terms $\lambda_k^e(t+t_l)$ and $\lambda_k^m(t+t_l)$ are the ratio of the variance explained to the unexplained variance of the extrapolation and NWP model components respectively. This combination is underpinned by the assumption that the three cascades are uncorrelated.

If there is no precipitation observed by the radar, then the stochastic forecast will initially forecast no rain, since the parameters controlling the noise cascade are taken from the radar cascade. As the forecast progresses these parameters will be merged with those determined from the NWP model. Therefore, this scheme can predict the structure and timing of precipitation not currently visible to the radars.

2.7 Merging of the Velocity Fields

For each time step it is necessary to advect the noise and extrapolation cascades, since the Lagrangian evolution of these cascades is being considered. This must be performed using a velocity field which represents the diagnosed advection from rain analyses at short lead times and the apparent velocity of the NWP forecasts precipitation at longer lead times. The apparent advection velocity of precipitation in NWP forecasts is estimated using the method described in section 3, using NWP forecasts in place of radar analyses. The merged advection velocity is calculated from

$$\mathbf{v}(t+t_l) = w_2^e(t+t_l)\mathbf{v}_e(t+t_l) + (1-w_2^e(t+t_l))\mathbf{v}_m(t+t_l) \quad (30)$$

where $\mathbf{v}_m$ is the velocity diagnosed from NWP forecasts and $\mathbf{v}_e$ is a perturbed velocity field derived from rain analyses (see section 3). The 2^{nd} level of the cascade, corresponding to spatial scales in the range 128 - 512 km, is used for the weights, since the climatological value of the auto-correlation coefficients for this level (see Table 2) most closely matches the measured auto-correlation of the velocity field.

The advection of the noise cascade ensures the correct spatio-temporal development of the precipitation pattern, but does pose a practical problem. Upon

advection of the cascade some data will be lost to numerical diffusion and some will be lost by advection out of the domain. Therefore, the cascade must be re-filtered to ensure that power is not lost at the lowest levels of the cascade. However, the Fourier filtering process cannot use the no-data areas which have been advected into the domain from the boundaries. These areas are replaced by Gaussian noise $N(0,1)$ before filtering. In this way the new data may be injected from the boundaries of the domain, and the noise cascade will always contain valid data at every point. However, the introduction of noise may affect the data values in the cascade in a non-local manner (at large-scales the cascade may be slightly adjusted over the whole domain). However, the new data is likely to be swamped by the data from the previous time step, and this should not pose a problem.

2.8 Re-constructing the Forecast

Finally, the output forecast field at time $t+t_l$ is calculated using

$$Z_{i,j}(t+t_l) = \sum_{k=1}^{n} \left(\mu_k(t+t_l) + \sigma_{k,i,j}(t+t_l) Y_{k,i,j}(t+t_l)\right). \tag{31}$$

The values of μ_k and $\sigma_{k,i,j}$ at time $t+t_l$ are a weighted combination of the values estimated from the latest rain analysis and NWP forecast (with variance adjusted). The estimates are

$$\mu_k(t+t_l) = \frac{w_k^e(t+t_l)}{w_k^e(t+t_l)+w_k^m(t+t_l)}\mu_k^e(t+t_l) + \frac{w_k^m(t+t_l)}{w_k^e(t+t_l)+w_k^m(t+t_l)}\mu_k^m(t+t_l) \tag{32}$$

and

$$\begin{aligned} \sigma_{k,i,j}(t+t_l) \quad = \quad & \frac{w_k^e(t+t_l)}{w_k^e(t+t_l)+w_k^m(t+t_l)}\sigma_{k,i,j}^e(t+t_l) + \\ & \frac{w_k^m(t+t_l)}{w_k^e(t+t_l)+w_k^m(t+t_l)}\sigma_{k,i,j}^m(t+t_l) \end{aligned} \tag{33}$$

where the symbols have their usual meaning and σ^e is advected using the merged velocity field.

The model formulation described above dictates that the forecast precipitation fields move gradually from an extrapolation nowcast to an NWP forecast. The scale-filtered noise ensures the forecasts are representative of the observed structure of precipitation. Since the temporal persistence of precipitation features is broadly proportional to their spatial extent, smaller scale features should decorrelate faster than larger scale features. Consequently, the noise component will rapidly dominate at the smallest scales and will then extend its influence upwards through the cascade to progressively larger scales as forecast lead time advances. This increasing influence will only be arrested by the introduction of NWP data.

3 Uncertainty due to Advection Errors

In the model formulation described here we have chosen to describe errors in the extrapolation forecast using a combination of uncertainties in the Lagrangian evolution of the rain rate field and uncertainties in the advection of that field. The

advection velocity is diagnosed using the optical flow method developed for the Gandolf nowcasting system [4]. The scheme has been enhanced slightly to improve its performance.

In weather scenarios with fast-moving areas of rain the optical flow scheme can suffer from problems with the numerical accuracy of the derivative estimation. This is avoided by applying a single advection vector to the entire domain before the optical flow scheme is applied. The advection vector is found by maximising the correlation between the latest radar analysis and the (displaced) previous analysis.

The Gandolf system [4] uses a forward-in-time advection scheme which allows separate rain areas to have independent velocities. However, the "banding" effects where stripes appear in areas where the velocity field is divergent were considered unacceptable in the probabilistic forecasting scheme. Thus, a backward-in-time advection scheme is used here. This requires that the minimisation algorithm (eq. 11 in [4]) is repeated with all blocks being given an equal weighting to ensure smooth variation of the velocity field across the entire domain.

Even with the large amount of spatial smoothing applied, the velocity field is subject to considerable noise. The diagnosed east-west velocity (usually denoted u) at a given point in the analysis (figure 4) shows that the estimated velocity appears to be a noisy quantity. The variogram for u, figure 5, shows the variance of the diagnosed velocity field which is

$$\sigma_u^2(t_d) = \langle u_{i,j}(t) u_{i,j}(t + t_d) \rangle \tag{34}$$

where $< . >$ denotes the average over all start times t and locations in the analysis i, j and the time lag is given by t_d. The value of the variance extrapolated to a time lag of zero gives the noise in the measurement of the velocity, and the value extrapolated to infinite time lag gives the variance of the true velocity field plus the measurement noise. Ideally, the measurement noise would be zero, but in this case it is larger than the variance in the true velocity field. Since the variance of the true velocity field is only apparent over a period of many hours we may safely introduce temporal smoothing of the diagnosed velocity to remove some of the measurement noise - without significantly affecting our ability to capture the variability intrinsic to the true velocity field. We choose to consider an exponential smoothing in time of the form

$$\mathbf{v}_{smooth}(t) = a\mathbf{v}_{smooth}(t - \Delta t) + (1 - a)\mathbf{v}_{estimate}(t) \tag{35}$$

where $\mathbf{v}_{smooth}$ is the temporally smoothed velocity, $\mathbf{v}_{estimate}$ is the unsmoothed estimate of the velocity, and a is a parameter giving the time-scale of the smoothing. Exponential smoothing in time is the natural choice since it is easy to implement and there is no sharp cut-off after which velocities are no longer used. If it was not possible to calculate a value of $\mathbf{v}_{estimate}$ at any time, then the smoothed velocity is not calculated. At subsequent times the last value of $\mathbf{v}_{smooth}$ that was calculated is used as the starting point for the current smooth velocity field. If the radar data is unavailable for more than two hours the estimation procedure is re-started. The mean square error of forecasts for June 2002 using different values of a is shown in figure 6. This shows that introducing temporal smoothing improves the skill of

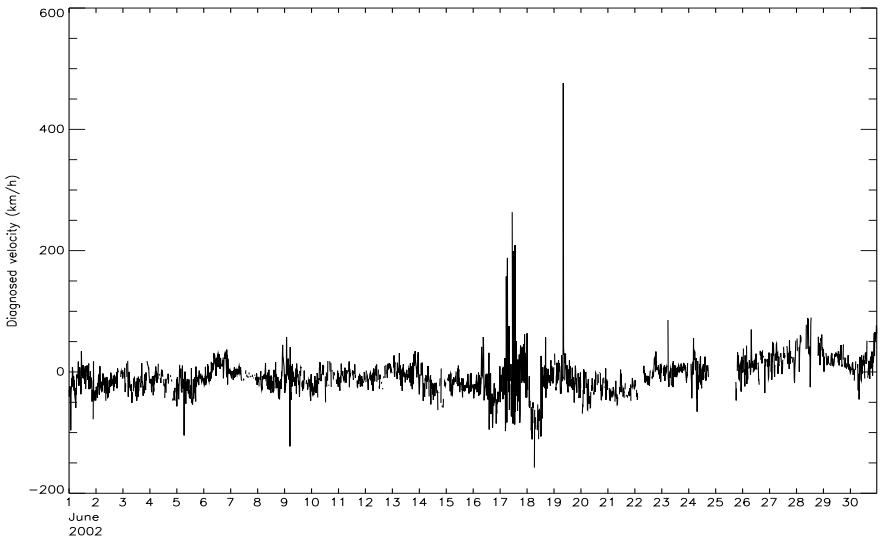

Figure 4: The diagnosed (unsmoothed) east-west velocity at a fixed point in the Gandolf domain. The velocity is not plotted when no data are available to diagnose the velocity. Note that this appears to be a noisy time series.

the forecast, but that too much smoothing produces a worse forecast. There is little to chose between the values of $a = 0.8515$ and $a = 0.5256$, their measured performance being very similar. Between these values of a the variability of the velocity field does not play a significant role, and the measurement noise has been effectively minimised. A value of $a = 0.85$ is chosen as appropriate, which gives a time-scale of the smoothing of approximately 90 minutes. The smoothed east-west velocity for the same point as in figure 4 is shown in figure 7.

A variogram for the smoothed velocities is shown in figure 8. The measurement noise is approximately zero, and the variability of the underlying process is unchanged from the unsmoothed case. Hence, the smoothing removes the noise in the measurement without affecting its ability to capture the evolution of the velocity field. Therefore, the evolution of the velocity field is slow and only changes over time-scales of many hours - a commonly held assumption. This assumption is unlikely to hold on small scales, but the broad scale motion which we are tracking is observed to change slowly.

We may estimate the significance of advection errors by comparing forecasts made using the smoothed velocity field with forecasts made using the smoothed velocity fields at future times. An advection forecast is normally made by advecting the observed rain rate field at T+0 by the diagnosed velocity field at T+0. To estimate the advection errors in a forecast one may use the diagnosed velocity field at T+15, T+30 and so on to advect the observed rain rate field at T+0. The difference in the mean square error (MSE) between these two forecasts should give an estimate of the importance of advection errors. From figure 9 we see

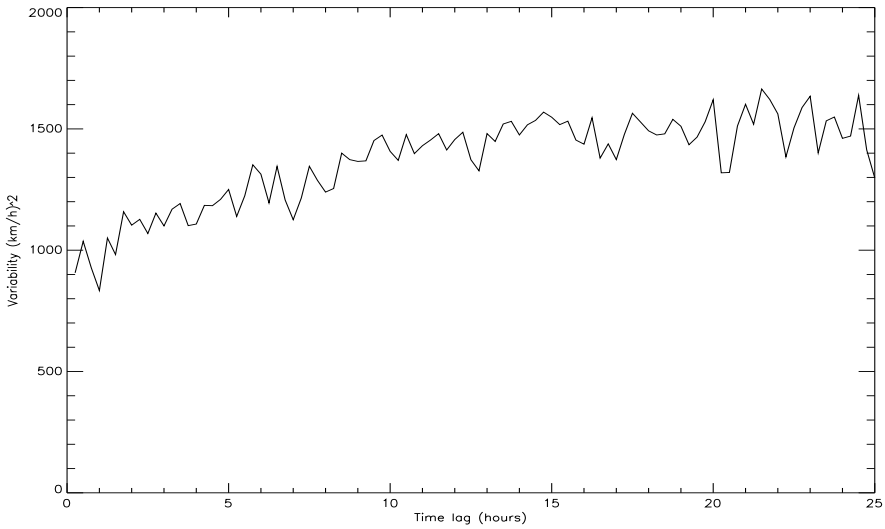

Figure 5: The variogram for the diagnosed (unsmoothed) east-west velocity (the data is similar for the north-south velocity). The plotted quantity is $\langle u_{i,j}(t)u_{i,j}(t+t_d)\rangle$ where $< . >$ denotes the average over all times t and locations in the forecast domain, t_d is the time lag and u is the east-west velocity. The measurement noise is greater than the variability in the apparent velocity. The variability of the velocity field is only significant over a period of many hours.

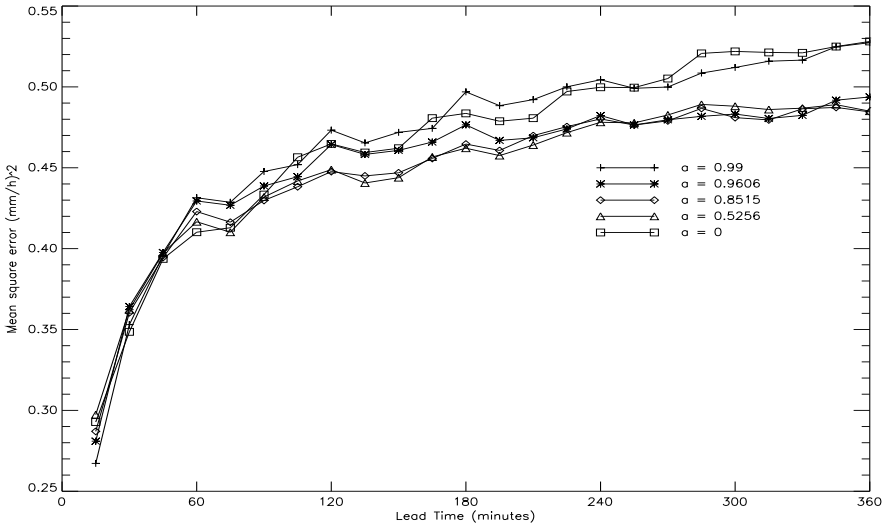

Figure 6: The mean square error for forecasts made with different levels of temporal smoothing, where a is used as given in equation 35. After a lead time of 2 hours it is clear that no smoothing and $a = 0.99$ are poor choices. A value of $a = 0.96$ is too high, since this performs poorly for lead times of less than 3 hours. A value of $a = 0.85$ was chosen since this seems to perform at least as well as any other value.

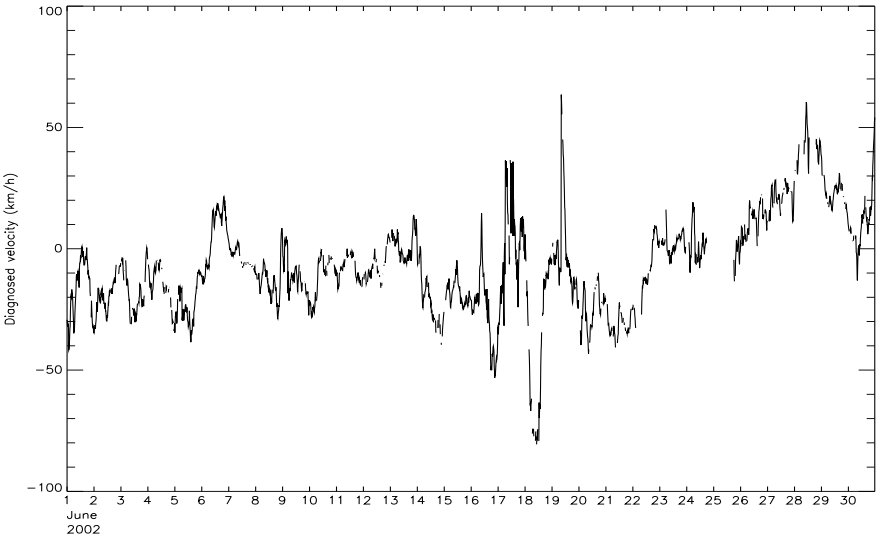

Figure 7: The smoothed east-west velocity at the same point as for figure 4. The velocity is not shown when no data are available to diagnose the velocity. Note that this appears to be a much less noisy time series than figure 4.

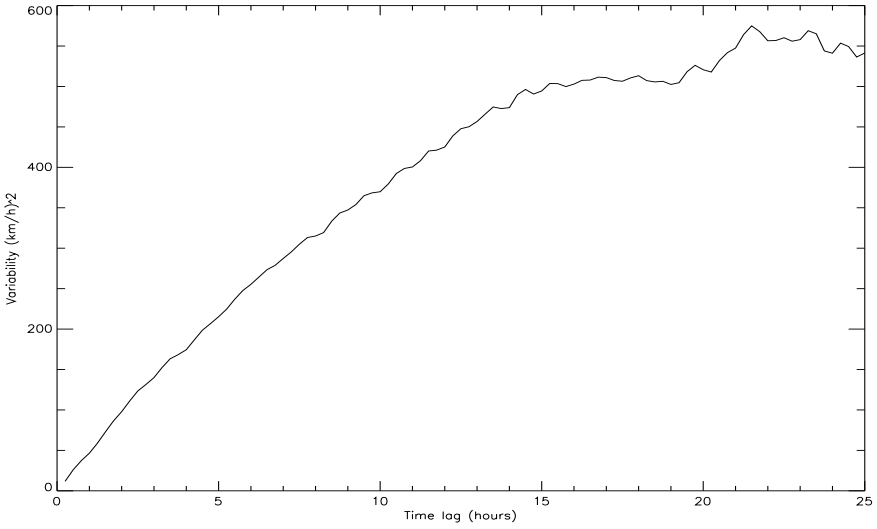

Figure 8: The variogram for the velocity field which has been temporally smoothed using $a = 0.85$. The measurement noise has been reduced to a very small value (indistinguishable from zero) without affecting the variability of the underlying process.

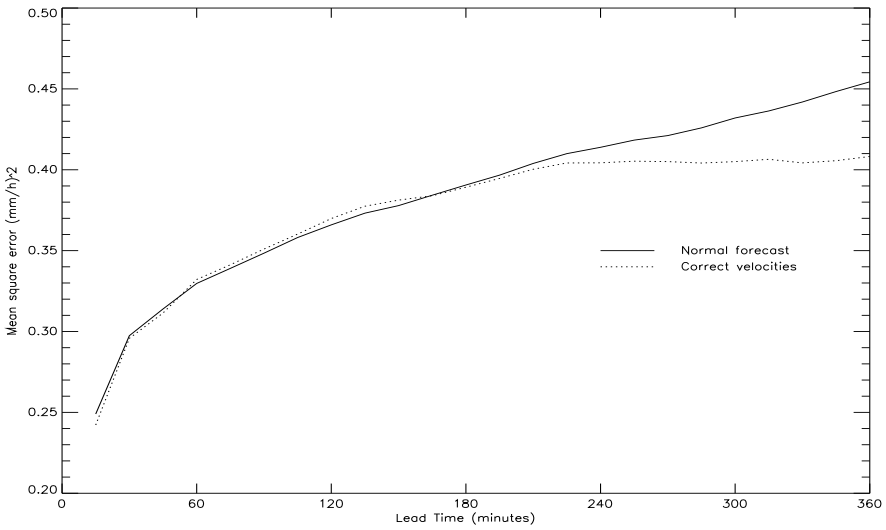

Figure 9: The mean square error of the forecasts using a smooth velocity field at the forecast time (normal forecast) and at future times (correct velocities). The correct velocities forecast is not significantly different from the normal forecast for lead times of less than three hours. This means that over these times the evolution of the apparent velocity field is insignificant. At a lead time of 6 hours the MSE of the correct velocities forecast is only 10% less than the normal forecast, indicating that the evolution of the velocity field has only a minor affect even at these lead times. These results may be dependent on the amount of spatial smoothing applied to the velocity field.

that advection errors are negligible for the first 3 hours of a forecast, but become significant by a lead time of 6 hours. Even then advection errors only account for around 10% of the total error in the forecast. By far the most significant source of error in the forecast is due to the evolution of the precipitation pattern. It should be noted at this point that the velocity fields produced from this advection scheme are highly smoothed in space, and these conclusions may be influenced by the level of spatial smoothing imposed.

To estimate the evolution of the velocity field we consider combinations of the differences $\mathbf{v}_{smooth}(t+t_d) - \mathbf{v}_{smooth}(t)$. From this it is apparent that the velocity field at time t is not an unbiased estimate of the velocity field at a future time $t+t_d$. Figure 10 shows the relation between the mean difference in the velocity field for a component parallel to the forecast velocity against the magnitude of the forecast velocity, that is $\langle(\mathbf{v}_{smooth}(t+t_d) - \mathbf{v}_{smooth}(t)).\hat{\mathbf{v}}_{smooth}(t)\rangle$ against $|\mathbf{v}_{smooth}(t)|$ where $\hat{\mathbf{v}}_{smooth}(t)$ denotes a unit vector in the direction of $\mathbf{v}_{smooth}(t)$ and $< ... >$ denotes the average over all points in the analysis and for all times in the forecast period of June 2002. This indicates that where the advection speed is diagnosed to be large it is unlikely to be as large (in that direction) in the future.

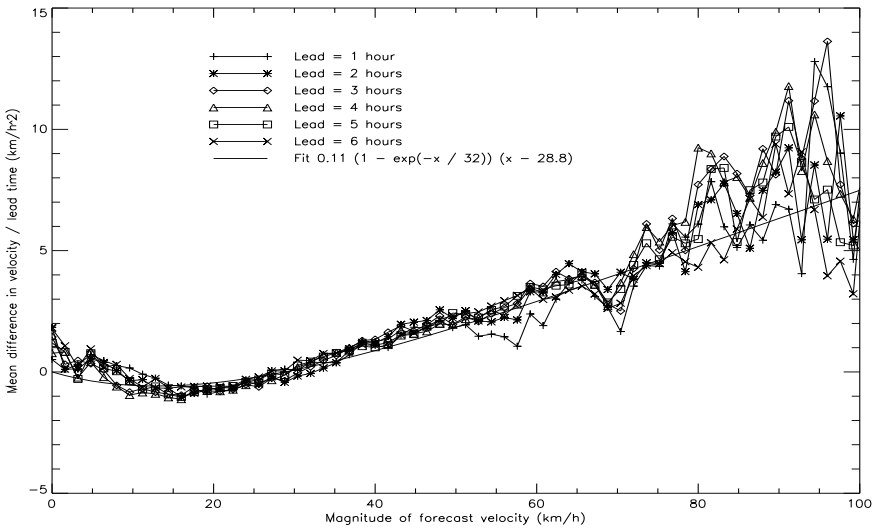

Figure 10: The bias of the forecast velocity in the direction parallel to forecast velocity ($\langle(\mathbf{v}_{smooth}(t+t_d) - \mathbf{v}_{smooth}(t)).\hat{\mathbf{v}}_{smooth}(t)\rangle$). This shows that the current diagnosed velocity is not an unbiased estimator of the future velocity, and must be reduced in magnitude to provide an unbiased estimate. This also helps to serve the role of reducing the variance of the original velocity field.

We believe that this is simply a regression of the future velocities towards the mean. Therefore, this forecast velocity will need to be reduced in magnitude before the velocity perturbation is added.

To check whether fast moving rain is more difficult to forecast than slow moving rain, we consider the correlation between the variance of the differences in the velocity field and the magnitude of the forecast velocity. Figure 11 shows $Var((\mathbf{v}_{smooth}(t+t_d) - \mathbf{v}_{smooth}(t)).\hat{\mathbf{v}}_{smooth}(t))$ where $Var(.)$ denotes the variance of the quantity in the argument. This does not appear to be correlated with the magnitude for the forecast velocity (noting to give much more weight to the left hand points on the graph since these velocities occur much more often), indicating that fast moving rain is *not* more difficult to forecast. Similar results are seen for the component of the differences perpendicular to the forecast velocity. Therefore, it is quite appropriate to add a velocity perturbation to the forecast velocity which is statistically homogeneous in space and uncorrelated with the forecast velocity. The following model will be used

$$\mathbf{v}_{noisy}(t+t_l) = \mathbf{v}_{smooth}(t)\left[1 - \frac{t_l}{60} f(|\mathbf{v}_{smooth}(t)|)\right] + \mathbf{v}_{perturbation}(t+t_l) \quad (36)$$

where t_l is the lead time of the forecast in minutes, $\mathbf{v}_{smooth}(t)$ is the diagnosed velocity, $f(x)$ is the bias removal / variance reduction described above which takes

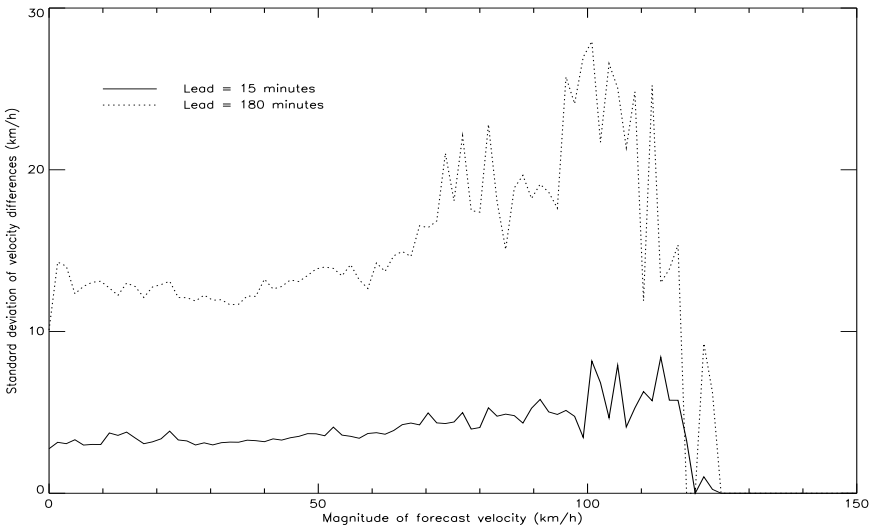

Figure 11: The variance of the differences in the velocity field, whose component is parallel to the forecast velocity. The variance does not increase (much) with increasing magnitude of the forecast velocity, indicating that fast moving areas of rain are *not* more difficult to forecast than slow moving areas of rain.

the form

$$f(x) = 0.11(x - 28.8)(1 - e^{-x/32}) \tag{37}$$

and $\mathbf{v}_{perturbation}(t_l)$ is the perturbation added to the velocity field.

Having removed the bias from the forecast velocity, reducing its magnitude by the appropriate factor, we may compare the forecast velocity with the future diagnosed velocity. The distribution of this quantity gives the distribution of the perturbation we should add to the advection velocity. For a lead time of 15 minutes, this distribution is very clearly a symmetric exponential distribution (figure 12). For longer lead times we consider the distribution of the following quantity

$$\mathbf{v}_{total}(t + t_l) = \mathbf{v}_{smooth}(t) \left[1 - \frac{(t_l + 1)}{120} f(|\mathbf{v}_{smooth}(t)|)\right] - \frac{1}{t_l} \sum_{i=1}^{t_l} \mathbf{v}_{smooth}(t + i) \tag{38}$$

which is the difference between the total forecast velocity and the total diagnosed velocity. This also follows a symmetric exponential distribution (see figure 13 for the distribution at a lead time of 120 minutes). The invariance of this distribution with lead time suggests a simple model for the uncertainty in the advection velocity - that the perturbation to add to the forecast velocity be a constant field for the whole forecast period. The variance of this distribution would vary with lead time, but the pattern could be maintained constant. Therefore we chose

$$\mathbf{v}_{perturbation}(t + t_l) = g(t_l)\mathbf{v}_{perturbation}(t) \tag{39}$$

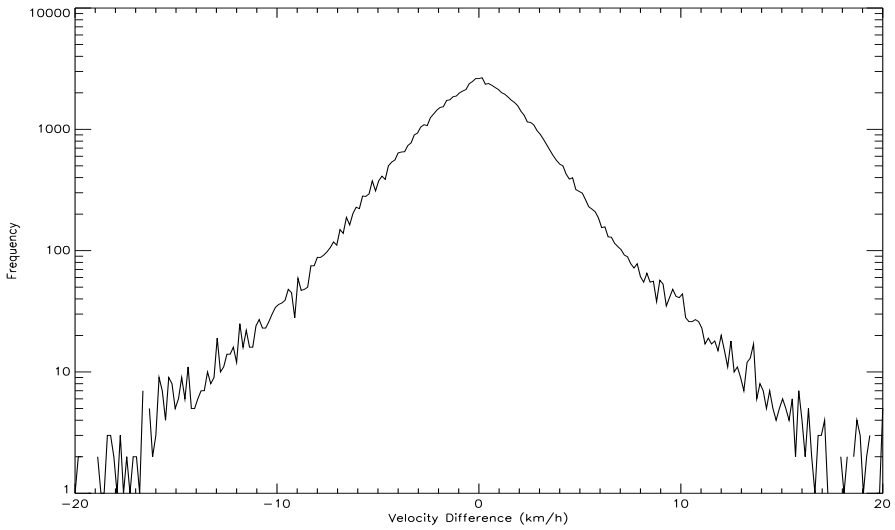

Figure 12: Distribution of the differences in velocity at 15 minutes lead time for the component parallel to the forecast velocity. This is an approximately symmetric exponential distribution.

where $g(t_l)$ is some function to be chosen, and $\mathbf{v}_{perturbation}(t)$ is the perturbation to the velocity field, taken as an exponential distribution with variance 1.

To determine the form of $g(t_l)$, figure 14 shows

$$\tilde{g}_{parallel}(t_l) = \sigma\left(\mathbf{v}_{total}(t+t_l).\hat{\mathbf{v}}_{smooth}(t)\right) - \sigma\left(\mathbf{v}_{total}(t+t_l-\Delta t).\hat{\mathbf{v}}_{smooth}(t)\right), \tag{40}$$

the differences of the standard deviation of the total differences in the velocity field, for components parallel and perpendicular to the forecast velocity. The following forms are chosen as appropriate fits

$$g_{parallel}(t_l) = -7.68 + 10.88 t_l^{0.23} \tag{41}$$

$$g_{perpendicular}(t_l) = -2.72 + 5.76 t_l^{0.31}. \tag{42}$$

We can apply this model to check that the distribution of the perturbed forecast velocities has the correct form (see figure 15). At a lead time of six hours, the distribution of the forecast velocities is slightly too wide, but this is considered sufficiently close for any difference to be ignored. The bias removal term $f\left(|\mathbf{v}_{smooth}(t)|\right)$ serves to reduce the variance of the deterministic part of the forecast velocity, but this is not quite sufficient to ensure the $\mathbf{v}_{noisy}$ has the correct distribution. This was calculated assuming that the smoothed diagnosed velocity has a Gaussian distribution. For June 2002, this was not the case (see figure 16), but we do not believe this discrepancy affects the results shown.

Finally, we need to ensure that the velocity perturbations $\mathbf{v}_{perturbation}(t)$ have the right spatial correlation structure, so that the forecast velocities appear correct.

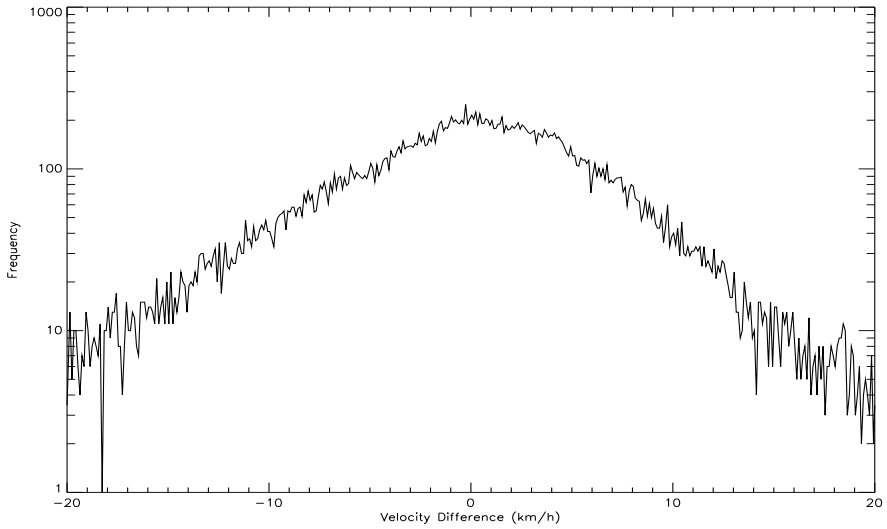

Figure 13: Distribution of the $\mathbf{v}_{total}$ for a forecast of 120 minutes lead time for the component parallel to the forecast velocity. This is an approximately symmetric exponential distribution. Since this follows the same distribution as the 15 minute differences, a simple way to generate the perturbations to the velocity is to add an unchanging field as the perturbation.

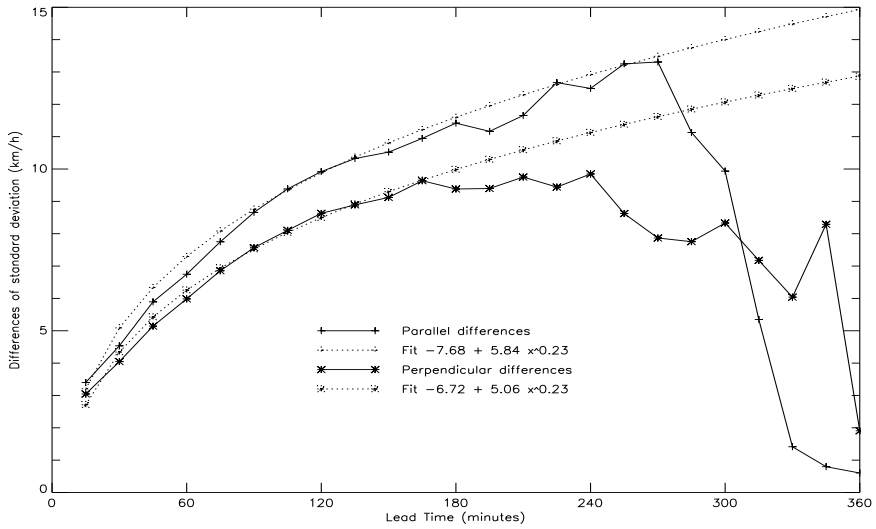

Figure 14: The standard deviation of the total differences in the forecast and diagnosed velocities, the component parallel and perpendicular to the forecast velocity.

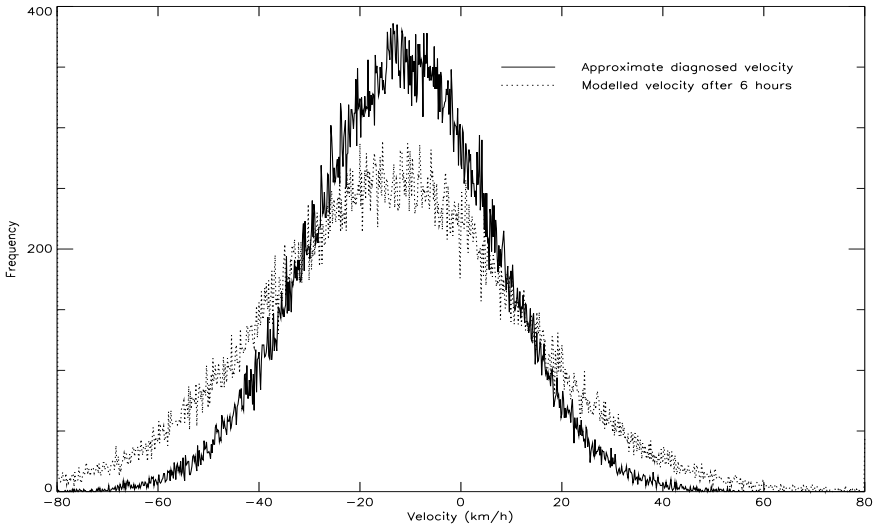

Figure 15: The modelled distribution of the perturbed east-west forecast velocity and the approximate distribution of the diagnosed velocity. We have assumed that the diagnosed velocity follows a Gaussian distribution and applied the model in equation 36 to generate the forecast velocity at a lead time of six hours. Ideally this would have the same distribution as the forecast velocity, but the discrepancy will be ignored at this stage.

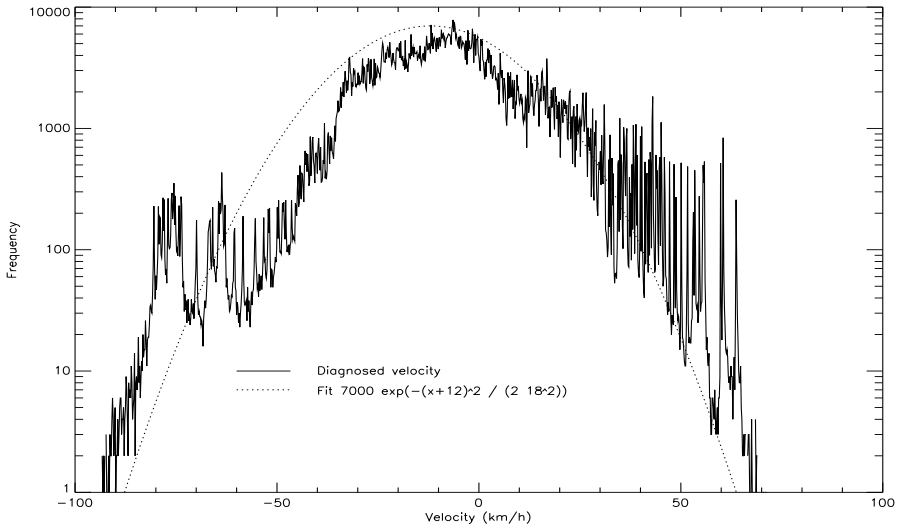

Figure 16: The distribution of the diagnosed east-west velocity - and a Gaussian fit to this distribution.

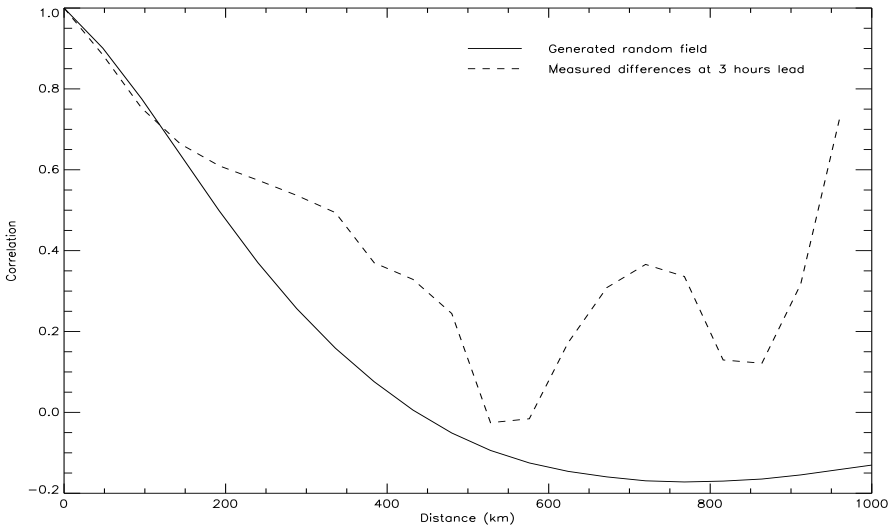

Figure 17: The spatial correlation of the total differences in the velocity for a forecast at a lead time of 3 hours, and that generated by the noise generation method. These agree well for the first 150 km.

This is achieved by generating a field of white noise (from a Gaussian distribution) and smoothing this field using the same minimisation of the Laplacian of the flow velocities used in the diagnosis program. This produces velocities which have the right correlation structure (see figure 17). When the velocity field is determined, the average motion over the whole domain is calculated first, before the optical flow method is applied. This is why there is a non-zero correlation over a separation of many 100 km, since noise in the rain analysis will not be diagnosed as having an advection velocity different from the field-wide average. This will not be accounted for by the perturbation scheme.

Since the velocity field diagnosed from the NWP model forecast is used at the same time at which it is used, this velocity field is not perturbed. Any uncertainty in the NWP forecast is taken as deriving from errors in the precipitation pattern.

4 Implementation Details

To evaluate the performance of the PQPN scheme described in sections 2 and 3 it has been run continuously for a period of ten days using Met Office meso-scale NWP precipitation forecasts and radar rain rate fields from the UK for March 2003. Data were ingested by the scheme every 15 minutes and an ensemble of forecasts produced every 3 hours to a maximum lead time of 6 hours. A forecast was run only when more than 5% of the rain analysis contained precipitation. If rain analyses were not available for more than 2 hours the estimation of the cascade parameters

and advection velocities was re-started afresh.

An ensemble of 100 members has been used to estimate the probability of precipitation at threshold rain rates of 0, 0.125, 0.25, 0.5, 1, 2, 4, 8 and 16 mm/h. For the calculation of the Brier scores the probabilities have been collected into 10% bins. For the ROC plots the probability thresholds of 1%, 2%, 5%, 10%, 20%, ... 90%, 96%, 99% and 100% have been used. Where the extrapolation forecast indicates a "no-data" value (such as outside the radar domain) the downscaled NWP forecast is used, otherwise the merging procedure described in section 2.6 is used. Since the perturbations added to each ensemble member are independent of all the other members the probability of precipitation forecast is generated by calculating the fraction of members which forecast precipitation to occur.

The computational requirements of the PQPN scheme are significantly greater than for existing, deterministic QPN schemes. To run for the required month took approximately 30 days of computer time on a Pentium 4 running at 2.4 GHz. Each single forecast of 100 ensemble members and one deterministic forecast during this time required slightly less than 8 hours of computer time. Therefore, a 50 member ensemble could be run every hour using 4 of these machines. The main portion of the computational expense is in re-filtering the noise cascade, and the Fourier transforms associated with that process. In the future, it may be possible to generate code that is more efficient. The overall flow of the program is shown in figure 18.

A best guess extrapolation forecast was run using components of our PQPN scheme for comparison with the probabilistic products. The former is very similar to the advection scheme implemented in the latest version of the Gandolf system [3, 4], but does not contain the convective life-cycle model and does not merge with NWP forecasts. The observed rain rate field is advected using a best estimate of the advection velocity field.

5 Results

5.1 Example Products

Figures 19 and 20 provides an example of deterministic (best guess advection) and probabilistic products with a data time of 03:00 GMT on 1 April 2003 for lead times of 2 and 6 hours. A front moved over the UK from the north east, followed by a region of showers. The NWP model forecast the broad pattern of the front and showers, but is unable to capture the finer structure of the precipitation pattern.

The probability of precipitation forecast shown in Figure 20 demonstrates how the forecast becomes less certain as the lead time increases. In forecast for a lead time of 1 hour (top-left) the join between the extrapolation forecast and the NWP forecast can still be seen. As the lead time increases and the influence of the NWP forecast within the area covered by the extrapolation forecast increases, this boundary becomes less distinct. It is interesting to note that the forecast indicates that no precipitation will occur in advance of the front, even at longer lead times.

Although example forecasts are interesting, any genuine benefits of the PQPN scheme over and above that afforded by deterministic schemes can only be identified

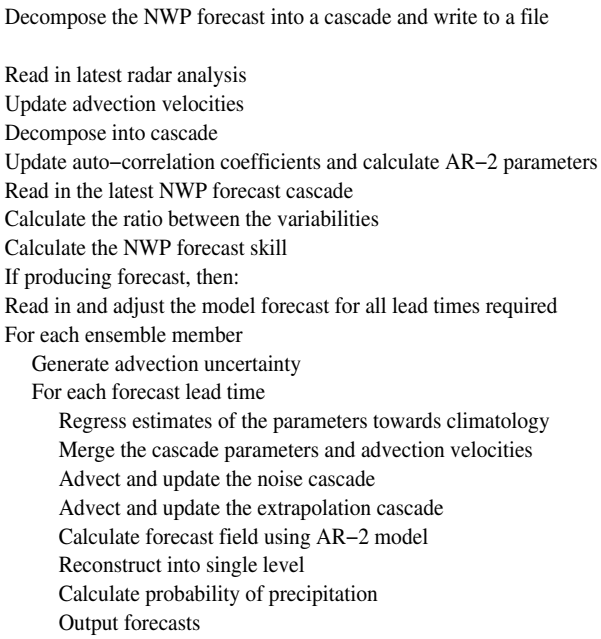

Figure 18: The flow of the forecast program. Each time an NWP forecast is produced, this is decomposed into a cascade, and this cascade stored in a file. This cascade is read in by the forecast program when it is required.

by a rigorous quantitative, statistical evaluation.

5.2 Brier skill score

The Brier skill score [6] is a useful measure for determining the skill of a probabilistic forecast. The Brier skill score is given as the Brier score of the forecast relative to the Brier score of a climatological forecast. It can be decomposed into 2 contributions [23] as follows

$$B_{skill} = B_{Resolution} - B_{Reliability} \tag{43}$$

where

$$B_{Resolution} = \frac{1}{Nf(1-f)} \sum_{c=1}^{N_c} F_c \left(\frac{O_c}{F_c} - f \right)^2 \tag{44}$$

measures the propensity of the forecast to give high or low values of the probability (as opposed to forecasting climatological values) and

$$B_{Reliability} = \frac{1}{Nf(1-f)} \sum_{c=1}^{N_c} F_c \left(\frac{O_c}{F_c} - p_c \right)^2 \tag{45}$$

is a penalty function for departure of the forecast from perfect reliability. f is the sample climatological value of the probability of precipitation, that is

$$f = \frac{\sum_{c=1}^{N_c} O_c}{\sum_{c=1}^{N_c} F_c}. \tag{46}$$

O_c is the number of observations of precipitation in the probability category c, F_c is the number of forecasts of precipitation in the same category, N_c is the number of categories, N is the total number of forecasts and p_c is the forecast probability of precipitation. Ideally, $p_c = \frac{O_c}{F_c}$ and the reliability penalty would be zero.

Figure 21 shows the reliability penalty for the forecast scheme. It is immediately clear that the forecast for high rain rate thresholds appears much less reliable than the forecast for low rain rate thresholds. It will of course be noted that the more intense rain rates are relatively rare, and so a probability forecast generated from our ensemble are less likely to capture the true probability of such "events" when compared with more common events of lower intensity. This may also indicate that there is a bias in the scheme, leading to an inaccurate forecast of the frequency of high rain-rate events.

To investigate the reliability issues further, Fig. 22 shows a reliability diagram for the rain / no-rain discrimination. This indicates that the forecast scheme tends to predict too much precipitation in the early stages of the forecast. As the forecast progresses this bias reduces. Therefore, there is need for more tuning of the scheme at short lead times to reduce the total precipitation area forecast.

The resolution (figure 23) and Brier skill score (figure 24) indicate that the scheme has predictive skill for lead times of up to six hours (and probably considerably beyond). The measured skill is greater at the lower rain rate thresholds because the occurrence of rain above these thresholds is more frequent, and any problems with reliability are only important at the highest rain rates.

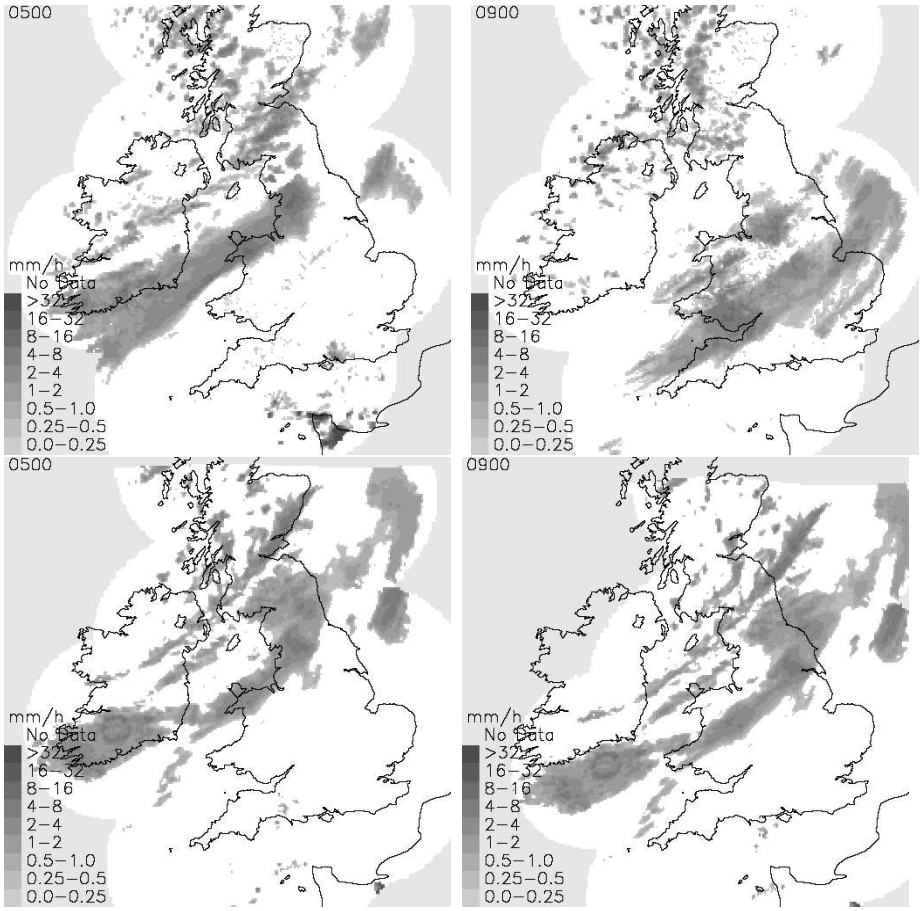

Figure 19: An example forecast made at 03:00 GMT on 1 April 2003. The top row shows the verifying analyses and the bottom row a deterministic extrapolation forecast. The left column is data valid at 05:00 GMT and the right column at 09:00 GMT.

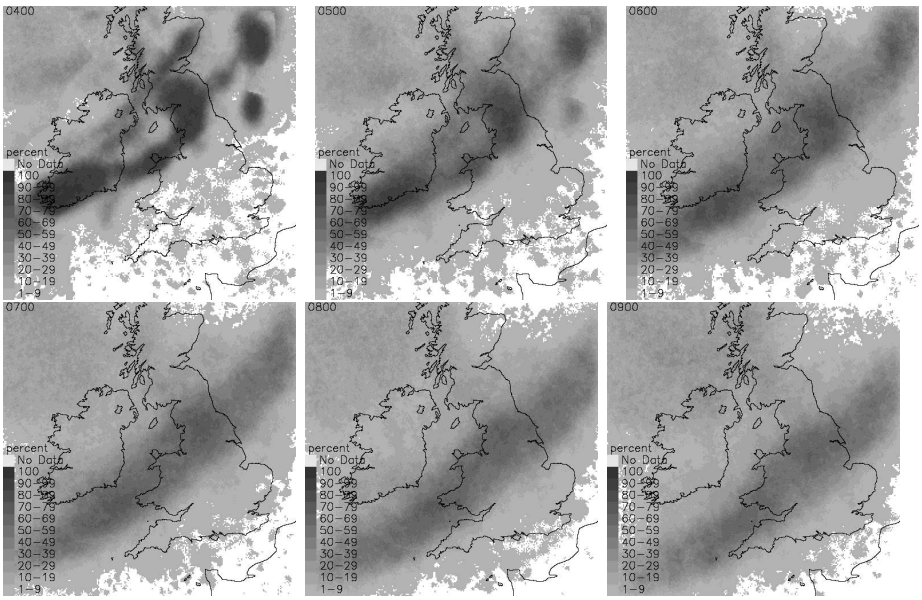

Figure 20: The probability of precipitation forecast made at 03:00 GMT on 1 April 2003. The forecast correctly diagnoses zero precipitation in advance of the front with an area of showers behind. From the longer lead time forecasts it is clear that only the very broad-scale pattern is taken from the NWP model. On the 04:00 forecast the boundary of the extrapolation forecast can be seen. As the lead time progresses and the NWP model forecast has greater influence, this boundary becomes less distinct

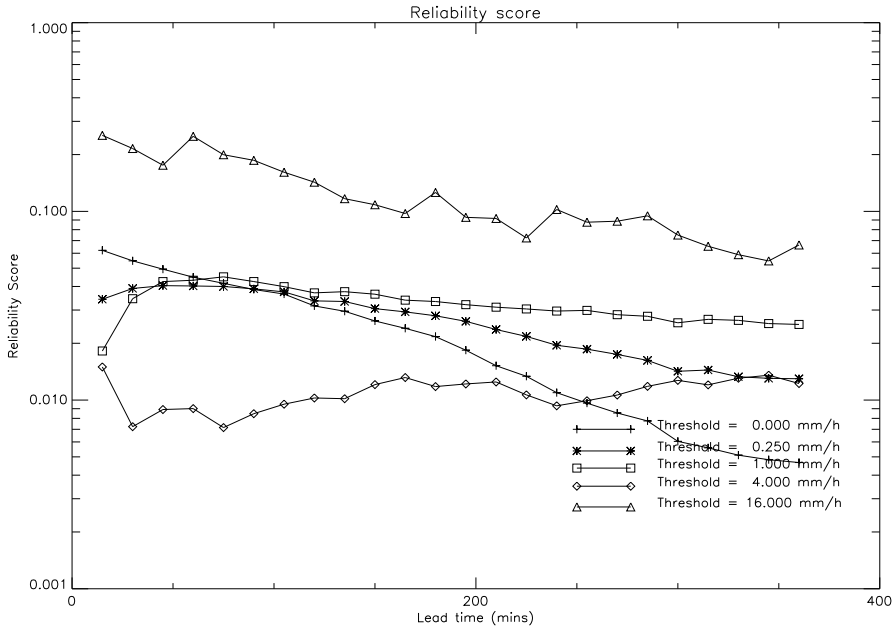

Figure 21: The reliability score - a penalty function for any departure from perfect reliability. The reliability appears to improve with increasing lead time, suggesting that STEPS has been better tuned for working with NWP forecasts than extrapolation forecasts.

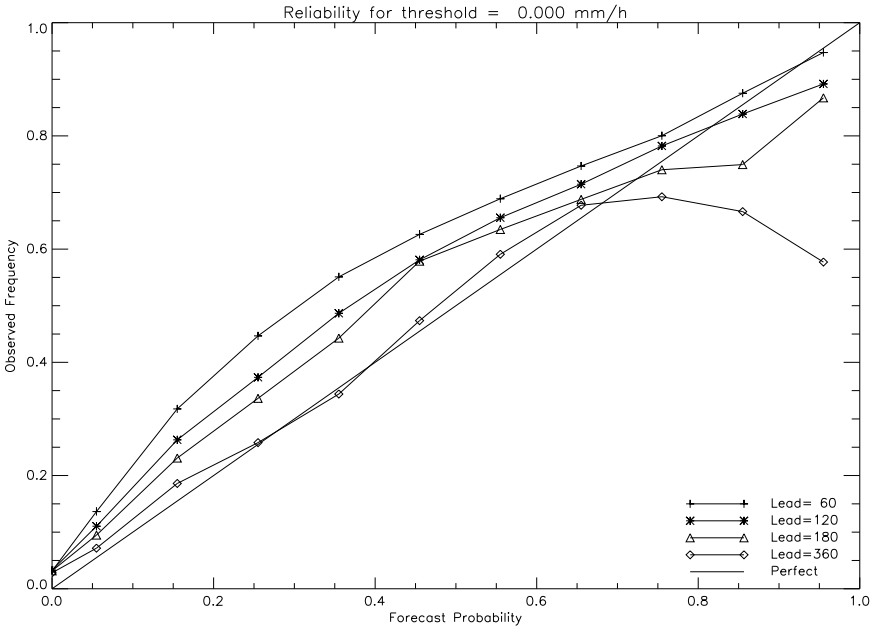

Figure 22: A reliability diagram for the rain / no-rain discrimination. The deviations from perfect reliability indicate that there is need for improved tuning of the forecasting scheme to reduce the amount of precipitation forecast in the early stages.

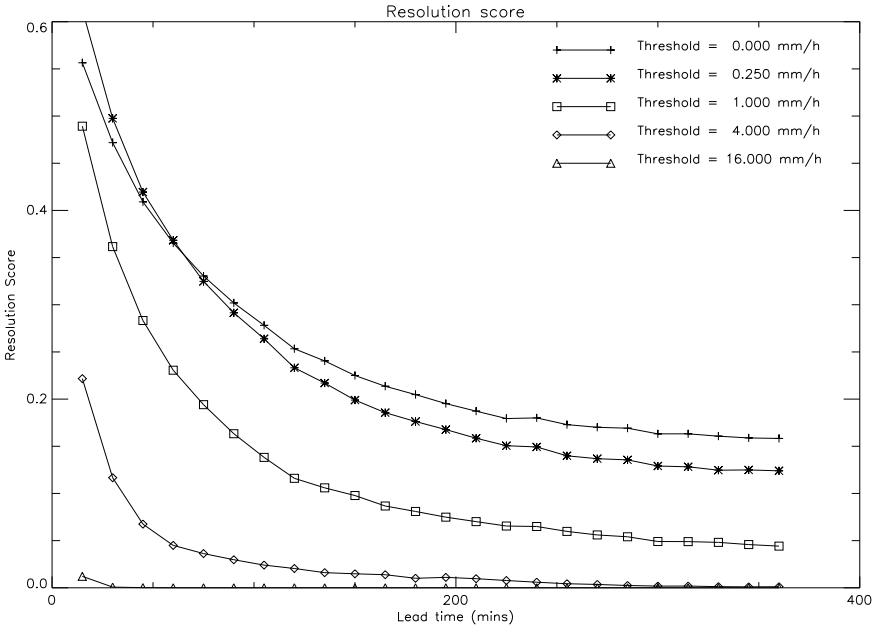

Figure 23: The resolution score which measures the difference between the forecast and climatology. The resolution decreases with increasing lead time, and is less for higher rain rates, as one might expect. The resolution at the highest rain rate threshold (16 mm/h) is essentially zero for all lead times.

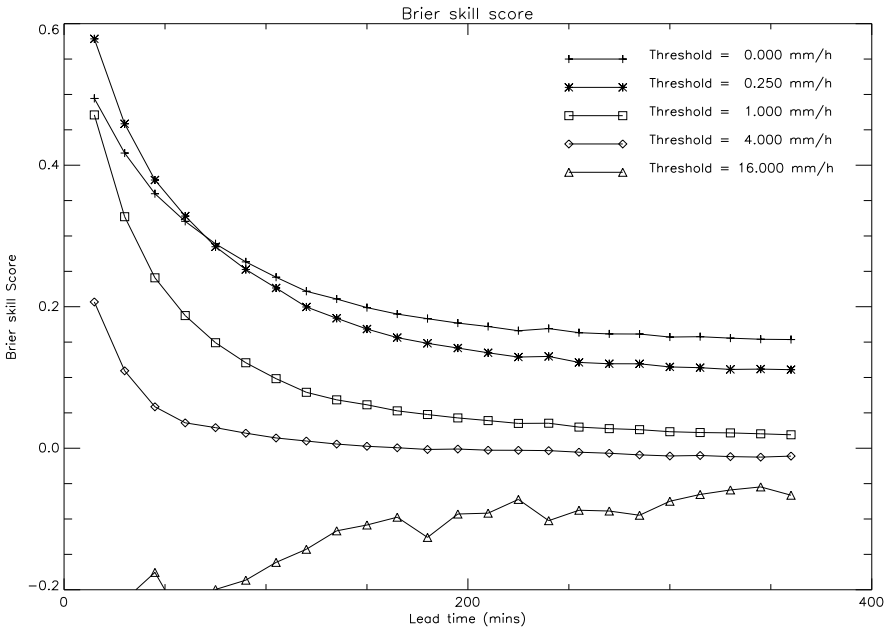

Figure 24: The Brier skill score - scores greater than zero indicate a skill over and above climatology. For the low rain rate thresholds the score indicates that the forecast has skill for over 6 hours and that the forecasts are reliable, since the skill score is dominated by the resolution term. For the highest rain rate threshold (> 16 mm/h) the skill is negative at all lead times. This skill score may be dominated by the base rate effect [35] so it is not necessarily meaningful to compare rain rate thresholds using this score.

5.3 Relative operating characteristic (ROC)

Plots of the relative operating characteristic [19] show the hit rate (H) against the false alarm rate (F) for different confidence levels (such as probability of precipitation greater than 50%). The hit rate and false alarm rate are defined as follows

$$H = \frac{a}{a+c} \tag{47}$$

$$F = \frac{b}{b+d} \tag{48}$$

where a, b, c and d are the standard contingency table values [35]. A measure of the skill of a PQPF is the area under the ROC curve (from (0,0) to (1,1)).

Figures 25 and 26 show the ROC for two different rain rate thresholds. For the rain / no-rain discrimination a hit rate of up to $\sim$ 85% is achieved for a lead time of 1 hour. However, this requires a false alarm rate slightly more than 30%. For a rain rate threshold of 2 mm/h the hit rate is never as high as for the rain / no-rain discrimination. Therefore the forecast at this threshold may seem to have less skill. Even at lead times of 6 hours the forecast has useful skill, showing the importance of including NWP model data (even if this is only used at the largest scales).

To achieve a high hit rate means that the forecast probability of precipitation must be non-zero over a wide area. For higher rain rate threshold this would mean forecasting a very low probability of precipitation over this wide area, since the frequency with which heavy rain is observed is low. To forecast a very low probability requires a very large ensemble - even with a threshold probability of 1% a large hit rate is not achieved. Hence forecasts for high rain rate thresholds will never achieve high hit rates, and therefore high skill scores. However, if the ROC curve for a fixed lead time and different thresholds is plotted, as in figure 27, one sees that the forecasts for the higher thresholds are almost as skillful as for the low thresholds. All the ROC curves fall onto the same line, indicating equal skill, even if high hit rates are unachievable at the higher thresholds due to a limited ensemble size.

5.4 Comparison with deterministic forecast

By plotting on an ROC curve the hit rate and false alarm rate for a deterministic forecast it is possible to compare a deterministic and stochastic forecast. Figure 25 shows this comparison for the rain / no-rain discrimination. Since the points for the deterministic forecast always lie below the respective curve for the stochastic forecast, we can conclude that the stochastic forecast is more skillful. This will be due to the beneficial influence of the NWP model forecast on the stochastic scheme. Unfortunately, we have been unable to run a deterministic forecast which merges the extrapolation and NWP components as this would provide a more meaningful test for the stochastic scheme.

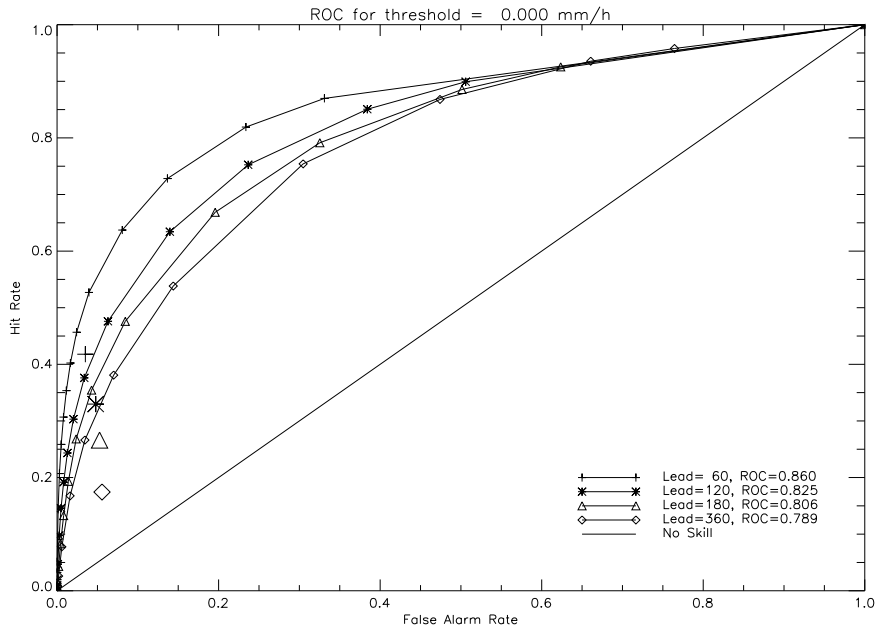

Figure 25: The ROC for the rain / no-rain discrimination. The forecast shows skill at all lead times. The legend shows the area under the curve, the ROC score. The probability thresholds used are (from top-right) 1%, 2%, 5%, 10%, 20% ... 90%, 96%, 99%, 100% on each curve. The same applies for Figures 26 and 27. The large symbols give the scores for the best-guess extrapolation forecast, which is seen to be consistently worse than the stochastic forecast.

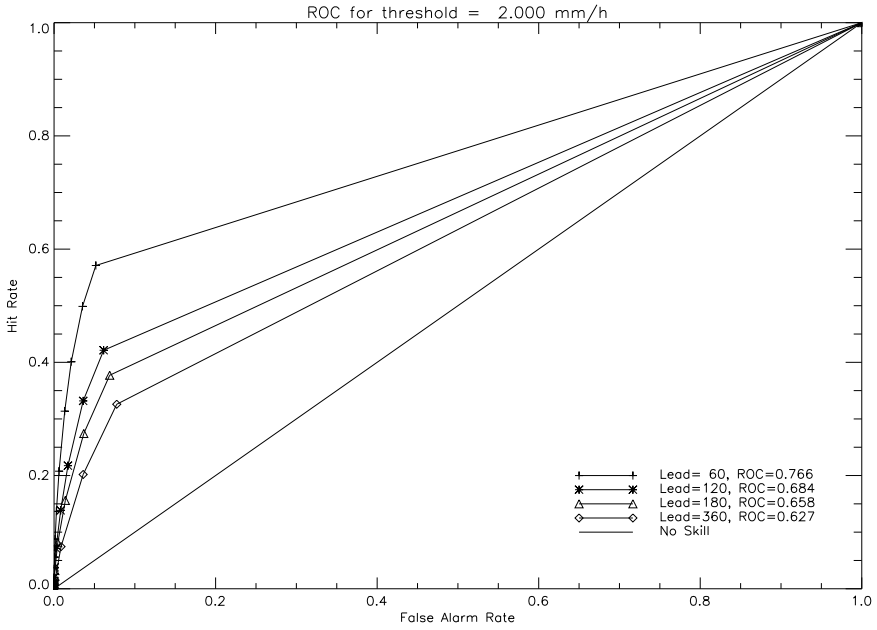

Figure 26: The ROC for a rain rate threshold of 2 mm/h. The forecast shows skill at all lead times, but is significantly diminished at T+6 hours. At this lead time the forecast is unable to achieve a hit rate better than $\sim 35\%$. The legend shows the area under the curve, the ROC score.

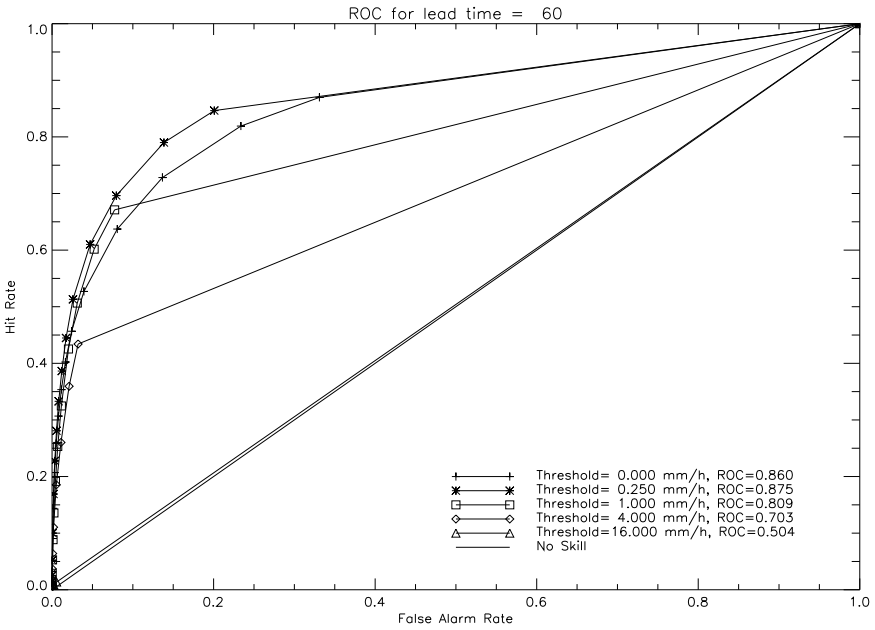

Figure 27: The ROC curve for forecasts made at a lead time of T+1 hour. This shows that the rain rate forecast for all the thresholds lie very closely together. The smaller value for the ROC at higher rain rate thresholds indicates that the scheme is unable to achieve high hit rates for these thresholds due to a limited ensemble size. One may wish to describe the scheme at high rain rates as skillful, even if forecasts of 1% probability (or less) of moderate to heavy precipitation are unlikely to be useful to many users.

6 Conclusions

A probabilistic precipitation nowcasting scheme which merges extrapolation nowcasts with downscaled NWP forecasts has been described. Both the uncertainty in the advection velocity and evolution of the rain rate distribution have been modelled. The scheme is therefore the first of its kind. The advection uncertainty is modelled by a random field of velocities which is correlated in space. The uncertainty in the field evolution is modelled using the S-PROG multiplicative cascade approach [32]. A hierarchy of AR-2 models on the levels in the cascade allows noise to be introduced most quickly at the smallest scales. An ensemble of forecasts is produced allowing the forecasts to be used by a number of other applications such as hydrological modelling.

This scheme produces reliable forecasts of the probability of precipitation at lead times up to six hours (although there is no necessary limit to the forecast period) and forecasting skill over climatology is seen for the entire forecasting period. Forecasting higher rain rates proved more difficult, but skill is seen out to six hours for thresholds up to 1 mm/h. Skill may exist with higher rain rates, but the probability of precipitation at any point is so low that it would require a vast ensemble size to accurately sample the PDF of the higher rain rates.

Since levels in the cascade corresponding to the smaller scales are affected by the larger scales (ringing) one may ask whether there is any meaning to the AR-2 model being applied to these levels? Any small change in the large scale features could have a large effect on the smaller scales in the cascade and the persistence of small scale features may be under-estimated. This effect has been minimised by the use of a Gaussian window in the filtering, but it is not possible to avoid this effect entirely.

The estimated rain rate contained within the rain analysis has been assumed to be error free since no credible estimates of the uncertainty in this product are available. It would be possible to include the uncertainty in the observations into this scheme if an ensemble of possible observations were available. Using such an ensemble would allow for the diagnosis of an ensemble of possible advection velocities and cascade parameters. Perhaps the most important requirement of such a product would be the correct characterisation of the spatial correlation of uncertainties in the rain analysis.

Great care has been taken to try to ensure that the forecast rain rate distribution closely mirrors the observed distribution of rain rates. However, intrinsic to the model development is that rain rates are log-normally distributed. This does not always appear to be the case, particularly for frontal passage where there may be a large area with moderate rain (around 1 to 4 mm/h) with only a few areas of light or heavy rain, as well as a large area with no rain. In such cases STEPS will tend to forecast too much light and heavy rain. A method for ameliorating this problem is to use a probability matching method [28] which will ensure the forecast rain rate distribution is always constant. Rather than ensuring that each ensemble member is matched with the latest rain analysis one should match the whole ensemble with the latest rain analysis. This avoids imposing any serious restriction on the ensemble and allows individual members to diverge.

The downscaling of the NWP forecast is based on the ratio between the variability in a radar cascade and the time-synchronous NWP forecast cascade estimated over a period of one month. This scheme is therefore not tuned specifically for dealing with extreme events, and an approach that adjusts the variability on the cascade using model diagnostics (such as total precipitable water or CAPE) may prove effective.

The variance of the rain / no-rain process is considerably larger than the variance of the raining process. This means that the variability in the lower cascade levels (small-scale features) is concentrated around the boundaries of the raining area. It is for this reason that the variability of the small scale features is allowed to vary spatially. It is also this difference that leads to the need to normalise the noise cascade to a variance different from unity. However, this adjustment is rather crude and does not perfectly re-create the structure of the input rainfall data.

Since the filtering applied to the data is isotropic, the features generated by the noise cascade tend to be circular in nature. This is not representative of atmospheric processes which are often aligned in bands. However, the extension to an anisotropic filter is likely to be prohibitively expensive.

A reliable forecasting scheme has been developed which has been subject to little tuning with most of the parameters of the model being estimated from the climatological values over a period of one month. Some work remains to ensure that the distribution of forecast rain rates is more precisely matched to the observations. It is hoped that with such work this scheme can be used operationally to aid with the production of fluvial flood warnings.

References

[1] T. Andersson and K-I. Ivarsson. A model for probability nowcasts of accumulated precipitation using radar. *Journal of Applied Meteology*, 30 No. 1:135–141, 1991.

[2] S. Applequist, G. E. Gahrs, E. L. Pfeffer, and X-F. Niu. Comparison of methodologies for probabilistic quantitative precipitation forecasting. *Weather and Forecasting*, 17 No. 4:783–799, 2002.

[3] N. E. Bowler and C. E. Pierce. Development of a short range qpf algorithm based on optic flow techniques. *Met Office NWP Technical Report No. 390*, 2002.

[4] N. E. Bowler, C. E. Pierce, and A. Seed. Development of a short range qpf algorithm based on optic flow techniques. In press - Journal of Hydrology, 2003.

[5] G. P. Box. *Time series analysis, forecasting and control*. Holden-Day, San-Fransisco, 1970.

[6] G. W. Brier. Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78:1–3, 1950.

[7] E. E. Ebert and J. L. McBride. Verification of precipitation in weather systems: determination of systematic errors. *Journal of Hydrology*, 239:179–202, 2000.

[8] E. Foufoula-Georgiou. On scaling theories of space-time rainfall, some recent results and open problems. In O. E. Barndorff-Nielsen, V. K. Gupta, V. Prez-Abreu, , and E. Waymire, editors, *Stochastic Methods in Hydrology: Rain, Landforms and Floods*, volume 7 of *Advanced Series on Statistical Science & Applied Probability*, pages 25–65. World Scientific, 1996.

[9] U. Germann and I. Zawadzki. Scale-dependence of the predictability of precipitation from continental radar images. part i: description of the methodology. *Monthly Weather Review*, 130 No. 12:2859–2873, 2002.

[10] B. W. Golding. Performance of the nimrod system. *Met Office Technical Report No. 226*, 1997.

[11] B. W. Golding. Nimrod: a system for generating automated very short range forecasts. *Meteorological Applications*, 5:1–16, 1998.

[12] D. Harris, E. Foufoula-Georgiou, K. K. Droegemeier, and J. J. Levit. Multiscale statistical properties of a high-resolution precipitation forecast. *Journal of Hydrometeorology*, 2:406–418, 2001.

[13] D. L. Harrison, S. J. Driscoll, and M. Kitchen. Improving precipitation estimates from weather radar using quality control and correction techniques. *Meteorological Applications*, 7 No. 2:135–144, 2000.

[14] P. Kumar and E. Foufoula-Georgiou. A multicomponent decomposition of spatial rainfall fields 1. segregation of large- and small-scale features using wavelet transforms. *Water Resources Research*, 29 No. 8:2515–2532, 1993.

[15] P. Kumar and E. Foufoula-Georgiou. A multicomponent decomposition of spatial rainfall fields 2. self-similarity in fluctuations. *Water Resources Research*, 29 No. 8:2533–2544, 1993.

[16] P. Kumar and E. Foufoula-Georgiou. Wavelet analysis for geophysical applications. *Reviews of Geophysics*, 35 No. 4:385–412, 1997.

[17] L. Li, W. Schmid, and J. Joss. Nowcasting of motion and growth of precipitation with radar over complex orography. *Journal of Applied Meteorology*, 34 No. 6:1286–1300, 1995.

[18] N. G. Mackay, R. E. Chandler, C. Onof, and H. S. Wheater. Disaggregation of spatial rainfall fields for hydrological modelling. *Hydrology and Earth System Sciences*, 5 No. 2:165–173, 2001.

[19] S. J. Mason and N. E. Graham. Conditional probabilities, relative operating characteristics, and relative operating levels. *Weather and Forecasting*, 14:713–725, 1999.

[20] S. Mecklenburg, J. Joss, and W. Schmid. Improving the nowcasting of precipitation in an alpine region with an enhanced radar echo tracking algorithm. *Journal of Hydrology*, 239:46–68, 2000.

[21] M. Menabde, A. W. Seed, D. Harris, and G. L. Austin. Multiaffine random field model of rainfall. *Water Resources Research*, 35 No. 2:509–514, 1999.

[22] F. Molteni, R. Buizza, T. N. Palmer, and T. Petroliagis. The ecmwf ensemble prediction system: methodology and validation. *Quarterly Journal of the Royal Meteorological Society*, 122:73–119, 1996.

[23] A. H. Murphy. A new vector partition of the probability score. *Journal of Applied Meteorology*, 12:595–600, 1973.

[24] G. G. S. Pegram and A. N. Clothier. Downscaling rainfields in space and time, using the string of beads model in time series model. *Hydrology and Earth System Sciences*, 5 No. 2:175–186, 2001.

[25] S. Percia and E. Foufoula-Georgiou. Linkage of scaling and thermodynamic parameters of rainfall: results from midlatitude mesoscale convective systems. *Journal of Geophysical Research*, 101 No. D3:7431–7448, 1996.

[26] S. Percia and E. Foufoula-Georgiou. Model for multiscale disaggregation of spatial rainfall based on coupling meteorological and scaling descriptions. *Journal of Geophysical Research*, 101 No. D21:26347–26361, 1996.

[27] C. E. Pierce and A. M. Cooper. Comparison of the performance of 2 km resolution object-oriented model and nimrod advection precipitation nowcast schemes. *Met Office Technical Report No. 350*, 2000.

[28] D. Rosenfeld, D. B. Wolff, and D. Atlas. General probability-matched relations between radar reflectivity and rain rate. *Journal of Applied Meteorology*, 32:50–72, 1993.

[29] J. D. Salas, J. W. Delleur, V. Yevevich, and W. L. Lane. *Applied modelling of hydrologic time series*. Water Resources Publications, 1980.

[30] W. Schmid, S. Mecklenburg, and J. Joss. Short-term risk forecasts of severe weather. *Physics and Chemistry of the Earth Part B*, 25 No. 10:1335–1338, 2000.

[31] W. Schmid, S. Mecklenburg, and J. Joss. Short-term risk forecasts of heavy rainfall. *Water Science and Technology*, 45 No. 2:121–125, 2002.

[32] A. W. Seed. A dynamic and spatial scaling approach to advection forecasting. *Journal of Applied Meteorology*, 42 No. 3:381–388, 2003.

[33] K. T. Smith and G. L. Austin. Nowcasting precipitation - a proposal for a way forward. *Journal of Hydrology*, 239:34–45, 2000.

[34] R. J. Standing. A report on a trial of probability forecasts for nimrod for possible application to the mobile phone short message service. (private communication), 1999.

[35] D. B. Stephenson. Use of the "odds ratio" for diagnosing forecast skill. *Weather and Forecasting*, 17 No. 2:221–232, 2000.

[36] Z. Toth and E. Kalnay. Ensemble forecasting at ncep and the breeding method. *Monthly Weather Review*, 125:3297–3319, 1997.

[37] D. Veneziano, R. L. Bras, and J. D. Niemann. Nonlinearity and self-similarity of rainfall in time and a stochastic model. *Journal of Geophysical Research*, 101:26371–26392, 1996.

A Recommendations for further work

The scheme described above is not ready for operational use, and has a few features which are undesirable. The following is an attempt to provide a list of these deficiencies and possible action to address them.

The scheme is computationally too expensive to be run at 2 km resolution, which is the standard required by the Environment Agency. Since the average auto-correlation coefficients at these scales are essentially zero (see Table 2), it would be possible to run the scheme at 4 km resolution, and quickly downscale the forecasts to 2 km resolution. The data required to transform between these resolutions corresponds to data with wavelengths between 2 and 4 km. Therefore, for each 4 km pixel, one chooses to define two randomly chosen numbers w_1, w_2 taken from a Gaussian distribution $N(0, \sigma_{n+1})$. The data values in the bottom level of the cascade would then be

$$\begin{matrix} w_1 & -w_2 \\ w_2 & -w_1 \end{matrix} \tag{49}$$

which may be added to all the other levels. The variance, σ_{n+1}, of the data at this level would need to be estimated off-line.

The climatological values of numerous parameters has been estimated using data from April 2003. However, if the model is run for a different season or at a different part of the globe, then these values may not be appropriate. It would be beneficial for the robustness of the code if the climatological values are updated automatically as the program runs.